

Oracle® Automatic Storage Management

Administrator's Guide

12c Release 2 (12.2)

E49628-07

July 2017

Oracle Automatic Storage Management Administrator's Guide, 12c Release 2 (12.2)

E49628-07

Copyright © 2007, 2017, Oracle and/or its affiliates. All rights reserved.

Primary Author: Eric Belden

Contributing Authors: Jerry Abraham, Prasad Bagal, Mark Bauer, Don Bolinger, Bill Bridge, Chan Chan, Xiaobin Chen, Hanlin Chien, Jonathan Creighton, Steve Fogel, Dave Friedman, Barb Glover, Fred Glover, Allan Graves, Cesar Gray, Shie-rei Huang, Rekha Iyer, Sameer Joshi, Ankur Kumar, Diane Lebel, Rich Long, Colin McGregor, Gregory Mengel, Harendra Mishra, Valarie Moore, Nitin Munjal, Darshan Nagarajappa, Harish Nandyala, Bob Nelson, Balaji Pagadala, Satish Panchumarthy, Hemanth Pikkili, Soma Prasad, Mounica Reddy, Kathy Rich, Charles Rusman, Harshavardhan Sabbineni, Francisco Sanchez, Santhosh Selvaraj, Ara Shakian, Duane Smith, Joshua Smith, Laura Solis, Malai Stalin, Amar Subba, Yaser Suleiman, Elisabeth Thibault, Michael Timpanaro-Perrotta, Samarjeet Tomar, Brian Tsao, Nitin Vengurlekar, P Venkatraman, Anton Verhulst, Radek Vingralek, Ron Weiss, Ric Werme, Rajiv Wickremesinghe, Jim A. Williams, Sivakumar Yarlagadda, Song Ye, Krishnan Yegnashankaran, Jon Zolla

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware 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 that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	xxxiii
Audience	xxxiii
Documentation Accessibility	xxxiii
Related Documents.....	xxxiv
Conventions.....	xxxiv
Changes in This Release for Oracle Automatic Storage Management	
Administrator's Guide	xxxv
Changes in Oracle Automatic Storage Management 12c Release 2 (12.2.0.1).....	xxxv
Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 2 (12.2.0.1)	xxxvii
Changes in Oracle Automatic Storage Management 12c Release 1 (12.1.0.2).....	xl
Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 1 (12.1.0.2)	xli
Changes in Oracle Automatic Storage Management 12c Release 1 (12.1.0.1).....	xlii
Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 1 (12.1.0.1)	xlvi
Part I Introduction	
1 Introducing Oracle Automatic Storage Management	
1.1 Overview of Oracle Automatic Storage Management	1-1
1.2 About Oracle Automatic Storage Management Cluster File System.....	1-2
1.3 Understanding Oracle ASM Concepts	1-3
1.3.1 About Oracle ASM Instances.....	1-3
1.3.2 About Oracle ASM Disk Groups.....	1-6
1.3.3 About Mirroring and Failure Groups.....	1-6
1.3.4 About Oracle ASM Disks	1-7
1.3.5 About Oracle ASM Allocation Units	1-7
1.3.6 About Oracle ASM Files	1-8
1.4 Understanding Oracle ASM Disk Group Administration.....	1-11
1.4.1 About Discovering Disks	1-11
1.4.2 About Mounting and Dismounting Disk Groups	1-12
1.4.3 About Adding and Dropping Disks	1-12

1.4.4 About Online Storage Reconfigurations and Dynamic Rebalancing	1-13
---	------

2 Exploring Considerations for Oracle ASM Storage

2.1 Storage Resources for Disk Groups.....	2-1
2.2 Oracle ASM and Multipathing	2-3
2.3 Recommendations for Storage Preparation.....	2-4
2.4 Storage Device Path and Permission Persistence.....	2-6

Part II Oracle ASM Instances and Disk Groups

3 Administering Oracle ASM Instances

3.1 Operating with Different Releases of Oracle ASM and Database Instances Simultaneously	3-1
3.2 Initialization Parameters for Oracle ASM Instances.....	3-2
3.2.1 About Initialization Parameter Files for an Oracle ASM Instance.....	3-3
3.2.2 About Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File.....	3-4
3.2.3 About Oracle ASM Initialization Parameter Settings	3-7
3.2.4 About Automatic Memory Management for Oracle ASM.....	3-7
3.2.5 Recommended Settings for Oracle ASM Initialization Parameters.....	3-8
3.2.6 About Setting Database Initialization Parameters for Use with Oracle ASM.....	3-16
3.3 Administration of Oracle ASM Instances.....	3-17
3.3.1 Using Oracle Grid Infrastructure for a Standalone Server.....	3-17
3.3.2 About Starting Up an Oracle ASM Instance.....	3-18
3.3.3 About Shutting Down an Oracle ASM Instance.....	3-22
3.3.4 About Administering Oracle ASM Instances with Server Control Utility	3-23
3.3.5 About Out of Place Upgrades.....	3-23
3.3.6 About Configuring Oracle Grid Infrastructure with the Configuration Wizard.....	3-24
3.3.7 About Active Session History Sampling for Oracle ASM.....	3-24
3.3.8 About Oracle Home User on Windows	3-24
3.4 Upgrading and Patching Oracle ASM.....	3-25
3.4.1 About Oracle ASM Rolling Upgrade.....	3-25
3.4.2 About Oracle ASM Rolling Patches.....	3-26
3.5 Authentication for Accessing Oracle ASM Instances.....	3-27
3.5.1 About Privileges for Oracle ASM.....	3-28
3.5.2 Creating Users with the SYSASM Privilege	3-31
3.5.3 Operating System Authentication for Oracle ASM.....	3-32
3.5.4 Password File Authentication for Oracle ASM.....	3-32
3.5.5 Managing a Shared Password File in a Disk Group	3-33
3.6 Administering the Oracle ASM Audit Trail	3-36
3.7 Migration of a Database to Use Oracle ASM.....	3-37
3.7.1 About Using Oracle Recovery Manager to Migrate Databases to Oracle ASM.....	3-37
3.7.2 Best Practices White Papers on Migrating to Oracle ASM.....	3-37

4 Administering Oracle ASM Disk Groups

4.1	Managing Disk Group Attributes	4-2
4.1.1	About Viewing and Setting Disk Group Attributes.....	4-2
4.1.2	Summary of Disk Group Attributes	4-3
4.2	Creating Disk Groups	4-9
4.2.1	Using the CREATE DISKGROUP SQL Statement.....	4-10
4.2.2	Creating Disk Groups for Data and Fast Recovery Area	4-12
4.2.3	Creating Disk Groups for a New Oracle Installation.....	4-14
4.2.4	Specifying the Logical and Physical Sector Size	4-15
4.2.5	Storing Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups	4-18
4.3	Altering Disk Groups	4-19
4.3.1	Managing Oracle ADVM Volumes in a Disk Group	4-21
4.3.2	Adding Disks to a Disk Group	4-22
4.3.3	Using the ALTER DISKGROUP SQL Statement to Add Disks to a Disk Group	4-22
4.3.4	Replacing Disks in Disk Groups	4-24
4.3.5	Renaming Disks in Disk Groups.....	4-24
4.3.6	Dropping Disks from Disk Groups.....	4-25
4.3.7	Intelligent Data Placement.....	4-26
4.3.8	Resizing Disks in Disk Groups.....	4-27
4.3.9	Undropping Disks in Disk Groups.....	4-28
4.3.10	Manually Rebalancing Disk Groups.....	4-28
4.3.11	Tuning Rebalance Operations	4-31
4.3.12	Scrubbing Disk Groups	4-32
4.4	Oracle ASM Disk Discovery.....	4-33
4.4.1	How A Disk is Discovered.....	4-33
4.4.2	Disk Discovery Rules	4-34
4.4.3	Improving Disk Discovery Time.....	4-35
4.5	Managing Capacity in Disk Groups.....	4-35
4.5.1	Negative Values of USABLE_FILE_MB.....	4-38
4.6	Oracle ASM Mirroring and Disk Group Redundancy	4-38
4.6.1	Mirroring, Redundancy, and Failure Group Options.....	4-38
4.6.2	Oracle ASM Recovery from Read and Write I/O Errors	4-42
4.6.3	Oracle ASM Fast Mirror Resync.....	4-43
4.6.4	Preferred Read Failure Groups	4-46
4.7	Performance and Scalability Considerations for Disk Groups	4-48
4.7.1	Determining the Number of Disk Groups.....	4-48
4.7.2	Performance Characteristics When Grouping Disks	4-48
4.7.3	Oracle ASM Storage Limits.....	4-49
4.8	Disk Group Compatibility	4-50
4.8.1	Overview of Disk Group Compatibility	4-50
4.8.2	Disk Group Compatibility Attributes	4-51
4.8.3	Setting Disk Group Compatibility Attributes	4-53

4.8.4	Features Enabled By Disk Group Compatibility Attribute Settings	4-57
4.8.5	Considerations When Setting Disk Group Compatibility	4-60
4.9	Managing Oracle ASM File Access Control for Disk Groups	4-61
4.9.1	About Oracle ASM File Access Control	4-62
4.9.2	Setting Disk Group Attributes for Oracle ASM File Access Control	4-63
4.9.3	Using SQL Statements to Manage Oracle ASM File Access Control	4-64
4.10	Maintaining Disk Groups	4-66
4.10.1	Mounting and Dismounting Disk Groups	4-67
4.10.2	Mounting Disk Groups Using the FORCE Option	4-67
4.10.3	Checking the Internal Consistency of Disk Group Metadata	4-68
4.10.4	Dropping Disk Groups	4-69
4.10.5	Renaming Disks Groups	4-70
4.11	Managing Oracle ASM Flex Disk Groups	4-72
4.11.1	About Oracle ASM Flex Disk Groups	4-73
4.11.2	About Oracle ASM File Groups	4-74
4.11.3	About Oracle ASM File Group Properties	4-76
4.11.4	About Quota Groups for Oracle ASM File Groups	4-79
4.11.5	About Oracle ASM Extended Disk Groups	4-80

5 Administering Oracle ASM Files, Directories, and Templates

5.1	What Types of Files Does Oracle ASM Support?	5-1
5.2	About Oracle ASM File Names	5-3
5.2.1	Single File Creation Form	5-4
5.2.2	Multiple File Creation Form	5-4
5.2.3	Fully Qualified File Name Form	5-5
5.2.4	Alias Oracle ASM File Name Forms	5-6
5.2.5	Creating a Tablespace in Oracle ASM: Using a Data File with an Alias Name	5-7
5.2.6	Alias Oracle ASM File Name with Template Form	5-7
5.2.7	Incomplete Oracle ASM File Name Form	5-8
5.2.8	Incomplete Oracle ASM File Name with Template Form	5-8
5.3	Creating and Referencing Oracle ASM Files in the Database	5-9
5.3.1	Creating Oracle ASM Files Using a Default File Location for Disk Group Specification	5-9
5.3.2	Using Oracle ASM File Names in SQL Statements	5-10
5.3.3	Moving Data Files Between Disk Groups With ALTER DATABASE	5-11
5.3.4	Oracle Database File Mapping for Oracle ASM Files	5-11
5.4	Managing Alias Names for Oracle ASM File Names	5-14
5.4.1	Adding an Alias Name for an Oracle ASM File Name	5-14
5.4.2	Renaming an Alias Name for an Oracle ASM File Name	5-14
5.4.3	Dropping an Alias Name for an Oracle ASM File Name	5-15
5.4.4	Dropping Files and Associated Aliases from a Disk Group	5-15
5.5	Managing Disk Group Directories	5-16
5.5.1	Creating a Directory	5-16

5.5.2	Renaming a Directory	5-16
5.5.3	Dropping a Directory	5-17
5.6	Accessing Oracle ASM Files with the XML DB Virtual Folder	5-17
5.6.1	Inside /sys/asm	5-18
5.7	Using DBMS_FILE Transfer Utility for Oracle ASM	5-19
5.8	Managing Disk Group Templates	5-19
5.8.1	Template Attributes	5-20
5.8.2	Adding Templates to a Disk Group	5-22
5.8.3	Modifying a Disk Group Template	5-23
5.8.4	Dropping Templates from a Disk Group	5-23
5.8.5	Creating Tablespaces in Oracle ASM: Specifying Attributes with Templates	5-23
6	Using Views to Display Oracle ASM Information	
6.1	Views Containing Oracle ASM Disk Group Information	6-1
6.2	Viewing Oracle ASM File Access Control Information	6-5
6.3	Viewing Disk Region Information	6-6
7	Administering Oracle ASM with Oracle Enterprise Manager	
7.1	Accessing the Oracle Automatic Storage Management Home Page	7-2
7.2	Configuring Oracle ASM Initialization Parameters with Oracle Enterprise Manager	7-2
7.3	Managing Oracle ASM Users with Oracle Enterprise Manager	7-3
7.4	Managing Disk Groups with Oracle Enterprise Manager	7-4
7.4.1	Creating Disk Groups	7-5
7.4.2	Adding Disks to Disk Groups	7-7
7.4.3	Bringing Disks Online and Offline	7-8
7.4.4	Dropping Disks from Disk Groups	7-9
7.4.5	Dropping Disk Groups	7-11
7.4.6	Monitoring Disk Group Usage	7-12
7.4.7	Mounting and Dismounting Disk Groups	7-12
7.4.8	Administering Advanced Disk Group Properties	7-13
7.4.9	Specifying a Power Limit For Disk Online	7-14
7.4.10	Replacing An Existing Disk With a New Disk	7-15
7.4.11	Setting and Modifying Content Type for the Disk Group	7-15
7.4.12	Setting and Modifying the Failure Group Repair Time	7-16
7.4.13	Running Rebalance on Multiple Disk Groups	7-16
7.4.14	Viewing Rebalance Process Related Data	7-17
7.4.15	Setting and Modifying the Thin Provisioned Disk Group Attribute	7-17
7.4.16	Running the On Demand Scrubbing Process	7-17
7.4.17	Determining Whether Scrubbing Process Is Running	7-18
7.5	Managing Oracle ASM File Access Control with Oracle Enterprise Manager	7-18
7.6	Managing Directories, Files, and Aliases with Oracle Enterprise Manager	7-19
7.7	Managing Disk Group Templates with Oracle Enterprise Manager	7-19
7.8	Managing Oracle ASM Filter Driver With Oracle Enterprise Manager	7-20

7.8.1	Discovering ASM Filter Driver Disks.....	7-20
7.8.2	Provisioning ASM Filter Driver Disks	7-20
7.8.3	Releasing ASM Filter Driver Disks	7-21
7.9	Monitoring Oracle ASM Performance With Oracle Enterprise Manager	7-21
7.9.1	Monitoring General Oracle ASM Performance.....	7-22
7.9.2	Checking Disk Group Integrity	7-22
7.10	Backing Up Oracle ASM Files with Oracle Enterprise Manager	7-22
7.11	Performing Bad Block Recovery with Oracle Enterprise Manager	7-23
7.12	Migrating to Oracle ASM with Oracle Enterprise Manager.....	7-23
7.13	Monitoring Alerts and Incidents With Oracle ASM Support Workbench.....	7-25
7.14	Monitoring Oracle Flex ASM	7-25
8	Performing Oracle ASM Data Migration with RMAN	
8.1	Overview of Oracle ASM Data Migration.....	8-1
8.1.1	Purpose of Oracle ASM Data Migration	8-2
8.1.2	Basic Concepts of Oracle ASM Data Migration	8-2
8.1.3	Basic Steps of Data Migration to Oracle ASM Using RMAN	8-3
8.2	Preparing to Migrate the Database to Oracle ASM Using RMAN	8-3
8.3	Migrating the Database to Oracle ASM Using RMAN	8-6
8.4	Migrating a Database from Oracle ASM to Alternative Storage	8-9
8.5	Moving Data Files Between Oracle ASM Disk Groups Using RMAN	8-10
9	Managing Oracle ASM With ASMCA	
9.1	Getting Started With the ASMCA GUI Tool.....	9-1
9.2	Managing Instances and Disk Groups With ASMCA.....	9-2
9.2.1	Creating Oracle ASM Instances with ASMCA.....	9-2
9.2.2	Managing Disk Groups with ASMCA	9-2
9.3	ASMCA Command-Line Interface.....	9-4
9.3.1	Running ASMCA Command-Line	9-4
9.3.2	ASMCA Commands for Instances and Disk Groups.....	9-5
10	Managing Oracle ASM with ASMCMD	
10.1	About ASMCMD	10-1
10.1.1	Types of ASMCMD Commands	10-2
10.1.2	About Oracle ASM Files, File Names, Directories, and Aliases	10-2
10.1.3	Preparing to Run ASMCMD.....	10-6
10.1.4	Running ASMCMD in Interactive Mode	10-7
10.1.5	Running ASMCMD in Noninteractive Mode	10-10
10.1.6	Getting Help.....	10-11
10.2	ASMCMD Instance Management Commands	10-11
10.2.1	dsget	10-13
10.2.2	dsset.....	10-13
10.2.3	lsct.....	10-14

10.2.4	lsop	10-15
10.2.5	lspwusr.....	10-15
10.2.6	orapwusr.....	10-16
10.2.7	pwdcopy.....	10-17
10.2.8	pwcreate.....	10-18
10.2.9	pwdelete.....	10-19
10.2.10	pwget.....	10-19
10.2.11	pwmove	10-20
10.2.12	pwset	10-21
10.2.13	showclustermode	10-22
10.2.14	showclusterstate	10-22
10.2.15	showpatches	10-22
10.2.16	showversion	10-23
10.2.17	shutdown.....	10-23
10.2.18	spbackup.....	10-24
10.2.19	spcopy	10-25
10.2.20	spget	10-27
10.2.21	spmove.....	10-27
10.2.22	spset.....	10-28
10.2.23	startup	10-29
10.3	ASMCMD Client Cluster Management Commands	10-30
10.3.1	lsc	10-30
10.3.2	mkcc.....	10-31
10.3.3	rmcc	10-32
10.4	ASMCMD Oracle ASM Filter Driver Management Commands	10-32
10.4.1	afd_configure	10-33
10.4.2	afd_deconfigure.....	10-34
10.4.3	afd_di.....	10-35
10.4.4	afd_dsget.....	10-36
10.4.5	afd_dsset.....	10-36
10.4.6	afd_filter.....	10-36
10.4.7	afd_label.....	10-37
10.4.8	afd_lsdsk.....	10-38
10.4.9	afd_lslbl.....	10-38
10.4.10	afd_refresh.....	10-39
10.4.11	afd_scan	10-40
10.4.12	afd_state	10-40
10.4.13	afd_unlabel.....	10-41
10.5	ASMCMD File Management Commands	10-42
10.5.1	amdu_extract.....	10-42
10.5.2	cd.....	10-43
10.5.3	cp.....	10-44
10.5.4	du.....	10-45

10.5.5	find.....	10-46
10.5.6	ls.....	10-47
10.5.7	lsof.....	10-49
10.5.8	mkalias.....	10-50
10.5.9	mkdir.....	10-51
10.5.10	pwd.....	10-51
10.5.11	rm.....	10-52
10.5.12	rmbalias.....	10-53
10.6	ASMCMD Disk Group Management Commands.....	10-53
10.6.1	chdg.....	10-54
10.6.2	chkdg.....	10-57
10.6.3	dropdg.....	10-57
10.6.4	iostat.....	10-58
10.6.5	lsattr.....	10-60
10.6.6	lsdg.....	10-61
10.6.7	lsdsk.....	10-63
10.6.8	lsod.....	10-65
10.6.9	md_backup.....	10-66
10.6.10	md_restore.....	10-67
10.6.11	mkdg.....	10-69
10.6.12	mount.....	10-71
10.6.13	offline.....	10-72
10.6.14	online.....	10-73
10.6.15	rebal.....	10-74
10.6.16	remap.....	10-75
10.6.17	setattr.....	10-75
10.6.18	stamp.....	10-76
10.6.19	stamplist.....	10-77
10.6.20	umount.....	10-78
10.7	ASMCMD File Group Management Commands.....	10-78
10.7.1	chfg.....	10-79
10.7.2	chqg.....	10-81
10.7.3	lsfg.....	10-81
10.7.4	lsqg.....	10-82
10.7.5	mkfg.....	10-83
10.7.6	mkqg.....	10-85
10.7.7	mvfg.....	10-85
10.7.8	rmfg.....	10-86
10.7.9	rmqg.....	10-86
10.8	ASMCMD Template Management Commands.....	10-87
10.8.1	chtpl.....	10-87
10.8.2	lstmpl.....	10-88
10.8.3	mktpl.....	10-89

10.8.4	rmtmpl	10-90
10.9	ASMCMD File Access Control Commands	10-90
10.9.1	chgrp.....	10-91
10.9.2	chmod.....	10-92
10.9.3	chown.....	10-93
10.9.4	groups	10-94
10.9.5	grpmod.....	10-95
10.9.6	lsgrp.....	10-96
10.9.7	lsusr	10-96
10.9.8	mkgrp.....	10-97
10.9.9	mkusr	10-98
10.9.10	passwd	10-98
10.9.11	rmgrp.....	10-99
10.9.12	rmusr	10-99
10.9.13	rpusr	10-100

Part III Oracle Automatic Storage Management Cluster File System

11 Introducing Oracle ACFS and Oracle ADVM

11.1	Overview of Oracle ACFS	11-1
11.2	Understanding Oracle ACFS Concepts	11-4
11.2.1	About Oracle ACFS.....	11-5
11.2.2	About the Oracle ACFS Mount Model and Namespace	11-5
11.2.3	About Oracle ACFS and Database Data Files	11-6
11.2.4	About Oracle ACFS and Oracle Database Homes.....	11-8
11.2.5	About Oracle ASM Dynamic Volume Manager	11-10
11.2.6	About the Oracle ACFS Driver Model	11-10
11.2.7	About the Oracle ACFS Mount Registry	11-11
11.2.8	About Oracle ACFS Snapshots.....	11-11
11.2.9	About Oracle ACFS and Backup and Restore.....	11-14
11.2.10	About Oracle ACFS Integration with Oracle ASM	11-15
11.2.11	About Oracle ACFS and External Tables on Windows	11-15
11.3	Understanding Oracle ACFS Administration	11-15
11.3.1	Oracle ACFS and File Access and Administration Security	11-16
11.3.2	Oracle ACFS and Grid Infrastructure Installation	11-17
11.3.3	Oracle ACFS Configuration	11-17
11.3.4	Oracle Clusterware Resources and Oracle ACFS Administration.....	11-18
11.3.5	Oracle ACFS and Dismount or Shutdown Operations.....	11-22
11.3.6	Oracle ACFS Security.....	11-22
11.3.7	Oracle ACFS Encryption	11-25
11.3.8	Oracle ACFS Compression.....	11-28
11.3.9	Oracle ACFS Auditing.....	11-29
11.3.10	Oracle ACFS Replication.....	11-36

11.3.11	Oracle ACFS Tagging	11-39
11.3.12	Using Replication with Auditing, Encryption, and Security	11-41
11.3.13	Oracle ACFS Plugins.....	11-43
11.3.14	Oracle ACFS Accelerator Volume.....	11-43
11.3.15	Oracle ACFS NAS Maximum Availability eXtensions	11-44
11.4	Overview of Oracle ASM Dynamic Volume Manager.....	11-49

12 Using Views to Display Oracle ACFS Information

12.1	Views Containing Oracle ACFS Information	12-1
12.2	Oracle ACFS Support for Oracle Database File Mapping Views	12-6

13 Administering Oracle ACFS with Oracle Enterprise Manager

13.1	Managing Oracle ACFS File Systems and Volumes with Oracle Enterprise Manager	13-2
13.1.1	Creating Oracle ACFS File Systems and Volumes	13-2
13.1.2	Viewing and Modifying Oracle ACFS Volumes and File Systems	13-3
13.2	Managing Oracle ACFS Snapshots with Oracle Enterprise Manager	13-4
13.2.1	Creating, Modifying, and Viewing Snapshots.....	13-4
13.2.2	Converting Snapshots Between Read-Only and Read-Write	13-5
13.2.3	Creating Child Snapshots from Existing Snapshots.....	13-5
13.3	Managing Security Features with Oracle Enterprise Manager	13-6
13.3.1	Initializing Oracle ACFS Security	13-7
13.3.2	Preparing an Oracle ACFS For Security.....	13-7
13.3.3	Undoing Oracle ACFS Security	13-7
13.3.4	Enabling Security on Oracle ACFS	13-8
13.3.5	Disabling Security on Oracle ACFS.....	13-8
13.3.6	Viewing Security Status.....	13-8
13.3.7	Adding an Oracle ACFS Security Administrator	13-9
13.3.8	Changing an Oracle ACFS Security Administrator Password	13-9
13.3.9	Deleting an Oracle ACFS Security Administrator.....	13-9
13.3.10	Viewing Oracle ACFS Security Administrators	13-9
13.3.11	Creating a Rule	13-10
13.3.12	Cloning a Rule.....	13-10
13.3.13	Editing a Rule.....	13-11
13.3.14	Deleting a Rule.....	13-11
13.3.15	Searching For and Viewing Rules.....	13-11
13.3.16	Creating a Ruleset	13-12
13.3.17	Cloning a Ruleset.....	13-12
13.3.18	Editing a Ruleset.....	13-13
13.3.19	Deleting a Ruleset.....	13-13
13.3.20	Searching For and Viewing Rulesets	13-13
13.3.21	Creating a Realm	13-14
13.3.22	Cloning a Realm	13-14
13.3.23	Adding Objects to a Realm	13-15

13.3.24	Deleting Objects from a Realm	13-15
13.3.25	Deleting a Realm.....	13-16
13.3.26	Viewing Realms	13-16
13.4	Managing Encryption Features with Oracle Enterprise Manager.....	13-16
13.4.1	Initializing Oracle ACFS Encryption	13-17
13.4.2	Enabling, Disabling, and Setting Parameters for Encryption On an Oracle ACFS	13-17
13.4.3	Viewing Encryption Status	13-17
13.5	Managing Tagging Features with Oracle Enterprise Manager.....	13-17
13.5.1	Adding a Tag in Oracle ACFS	13-18
13.5.2	Removing a Tag in Oracle ACFS.....	13-18
13.5.3	Searching for Tags in Oracle ACFS.....	13-19
14	Managing Oracle ACFS and Oracle ADVM With ASMCA	
14.1	ASMCA GUI Tool for Managing Oracle ACFS and Oracle ADVM.....	14-1
14.1.1	Managing Oracle ADVM Volumes with ASMCA.....	14-2
14.1.2	Managing Oracle ACFS File Systems with ASMCA	14-3
14.1.3	Managing Security and Encryption for Oracle ACFS with ASMCA	14-4
14.1.4	Creating an Oracle ACFS File System for Database Use	14-5
14.2	ASMCA Command-Line Interface for Managing Oracle ACFS and Oracle ADVM.....	14-6
14.2.1	ASMCA Commands for Oracle ACFS and Oracle ADVM	14-6
15	Managing Oracle ADVM with ASMCMD	
15.1	ASMCMD Volume Management Commands.....	15-1
15.1.1	volcreate.....	15-2
15.1.2	voldelete.....	15-4
15.1.3	voldisable.....	15-5
15.1.4	volenable.....	15-6
15.1.5	volinfo	15-6
15.1.6	volresize	15-7
15.1.7	volset	15-8
15.1.8	volstat	15-9
16	Managing Oracle ACFS with Command-Line Tools	
16.1	Basic Steps to Manage Oracle ACFS Systems.....	16-2
16.1.1	About Using Oracle ACFS Command-Line Tools.....	16-2
16.1.2	Creating an Oracle ACFS File System	16-5
16.1.3	Accessing an Oracle ACFS File System on a Different Node in the Cluster	16-7
16.1.4	Managing Oracle ACFS Snapshots	16-8
16.1.5	Securing Oracle ACFS File Systems.....	16-9
16.1.6	Encrypting Oracle ACFS File Systems	16-14
16.1.7	Tagging Oracle ACFS File Systems	16-16
16.1.8	Replicating Oracle ACFS File Systems	16-17

16.1.9	Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems.....	16-23
16.1.10	Removing an Oracle ACFS File System and a Volume	16-24
16.2	Oracle ACFS Command-Line Tools for Linux Environments	16-25
16.2.1	fsck.....	16-26
16.2.2	mkfs	16-28
16.2.3	mount	16-30
16.2.4	umount.....	16-32
16.3	Oracle ACFS Command-Line Tools for the Solaris Environment.....	16-33
16.3.1	fsck.....	16-34
16.3.2	mkfs	16-36
16.3.3	mount	16-38
16.3.4	umount/umountall.....	16-40
16.4	Oracle ACFS Command-Line Tools for the AIX Environment.....	16-41
16.4.1	fsck.....	16-42
16.4.2	mkfs	16-44
16.4.3	mount	16-46
16.4.4	umount/umountall.....	16-48
16.5	Oracle ACFS Command-Line Tools for Windows Environments	16-49
16.5.1	acfschkdisk	16-50
16.5.2	acfsdismount	16-52
16.5.3	acfsformat	16-53
16.5.4	acfsmountvol.....	16-54
16.5.5	acfsutil detach	16-56
16.5.6	advmutil dismount.....	16-56
16.5.7	advmutil list	16-58
16.5.8	advmutil mount.....	16-58
16.6	Oracle ACFS Command-Line Tools for Tagging	16-60
16.6.1	acfsutil tag info.....	16-60
16.6.2	acfsutil tag set.....	16-61
16.6.3	acfsutil tag unset.....	16-63
16.7	Oracle ACFS Command-Line Tools for Replication.....	16-63
16.7.1	acfsutil repl bg.....	16-64
16.7.2	acfsutil repl compare.....	16-65
16.7.3	acfsutil repl info	16-67
16.7.4	acfsutil repl init	16-69
16.7.5	acfsutil repl pause.....	16-72
16.7.6	acfsutil repl resume	16-73
16.7.7	acfsutil repl sync	16-73
16.7.8	acfsutil repl terminate	16-75
16.7.9	acfsutil repl trace	16-76
16.7.10	acfsutil repl update.....	16-76
16.7.11	acfsutil repl upgrade	16-78

16.8	Oracle ACFS Command-Line Tools for Auditing	16-81
16.8.1	acfsutil audit archive.....	16-81
16.8.2	acfsutil audit disable	16-82
16.8.3	acfsutil audit enable	16-82
16.8.4	acfsutil audit info.....	16-83
16.8.5	acfsutil audit init.....	16-84
16.8.6	acfsutil audit purge	16-85
16.8.7	acfsutil audit read.....	16-85
16.9	Oracle ACFS Command-Line Tools for Security	16-86
16.9.1	acfsutil sec admin add	16-87
16.9.2	acfsutil sec admin password.....	16-88
16.9.3	acfsutil sec admin remove.....	16-88
16.9.4	acfsutil sec batch	16-89
16.9.5	acfsutil sec disable	16-90
16.9.6	acfsutil sec enable	16-91
16.9.7	acfsutil sec info.....	16-91
16.9.8	acfsutil sec info file	16-92
16.9.9	acfsutil sec init.....	16-93
16.9.10	acfsutil sec load	16-95
16.9.11	acfsutil sec prepare	16-96
16.9.12	acfsutil sec realm add.....	16-98
16.9.13	acfsutil sec realm audit disable.....	16-102
16.9.14	acfsutil sec realm audit enable.....	16-103
16.9.15	acfsutil sec realm audit info	16-105
16.9.16	acfsutil sec realm clone	16-105
16.9.17	acfsutil sec realm create	16-107
16.9.18	acfsutil sec realm delete.....	16-108
16.9.19	acfsutil sec realm destroy	16-109
16.9.20	acfsutil sec rule clone	16-109
16.9.21	acfsutil sec rule create	16-110
16.9.22	acfsutil sec rule destroy	16-112
16.9.23	acfsutil sec rule edit.....	16-112
16.9.24	acfsutil sec ruleset clone	16-113
16.9.25	acfsutil sec ruleset create	16-114
16.9.26	acfsutil sec ruleset destroy	16-115
16.9.27	acfsutil sec ruleset edit.....	16-116
16.9.28	acfsutil sec save.....	16-116
16.10	Oracle ACFS Command-Line Tools for Encryption.....	16-117
16.10.1	acfsutil encr info	16-118
16.10.2	acfsutil encr init	16-119
16.10.3	acfsutil encr off.....	16-120
16.10.4	acfsutil encr on.....	16-121
16.10.5	acfsutil encr rekey.....	16-122

16.10.6	acfsutil encr set.....	16-123
16.10.7	acfsutil keystore migrate	16-125
16.11	Oracle ACFS Command-Line Tools for Snapshots.....	16-125
16.11.1	acfsutil snap convert	16-126
16.11.2	acfsutil snap create	16-127
16.11.3	acfsutil snap delete	16-130
16.11.4	acfsutil snap duplicate apply	16-131
16.11.5	acfsutil snap duplicate create.....	16-132
16.11.6	acfsutil snap info.....	16-134
16.11.7	acfsutil snap quota	16-136
16.11.8	acfsutil snap remaster	16-138
16.11.9	acfsutil snap rename	16-139
16.12	Oracle ACFS Command-Line Tools for Compression	16-140
16.12.1	acfsutil compress copy.....	16-140
16.12.2	acfsutil compress info	16-141
16.12.3	acfsutil compress off	16-142
16.12.4	acfsutil compress on.....	16-142
16.13	Oracle ACFS Command-Line Utilities	16-143
16.13.1	acfsutil accel replace.....	16-143
16.13.2	acfsutil compat get	16-144
16.13.3	acfsutil compat set	16-145
16.13.4	acfsutil defrag dir	16-146
16.13.5	acfsutil defrag file.....	16-146
16.13.6	acfsutil info file	16-147
16.13.7	acfsutil info fs.....	16-149
16.13.8	acfsutil info id	16-154
16.13.9	acfsutil plugin disable.....	16-155
16.13.10	acfsutil plugin enable.....	16-156
16.13.11	acfsutil plugin info	16-159
16.13.12	acfsutil registry	16-160
16.13.13	acfsutil rmfs.....	16-162
16.13.14	acfsutil scrub	16-163
16.13.15	acfsutil size	16-164
16.13.16	acfsutil tune	16-167
16.13.17	advmutil canonical.....	16-168
16.13.18	advmutil tune.....	16-169
16.13.19	advmutil volinfo.....	16-170

Part IV Advanced Topics

17 Understanding Oracle ACFS Advanced Topics

17.1	Limits of Oracle ACFS.....	17-1
17.1.1	Oracle ACFS Disk Space Usage.....	17-2

17.1.2	Oracle ACFS Error Handling.....	17-3
17.1.3	Oracle ACFS and NFS.....	17-3
17.2	Limits of Oracle ADVM.....	17-4
17.3	Oracle ACFS Loopback Support.....	17-4
17.4	Oracle ACFS Drivers Resource Management.....	17-4
17.5	Oracle ACFS Registry Resource Management.....	17-5
17.6	Oracle ACFS File System Resource Management.....	17-5
17.7	Oracle ACFS and Oracle Restart.....	17-6
17.8	Oracle ACFS Driver Commands.....	17-6
17.8.1	acfsload.....	17-6
17.8.2	acfsdriverstate.....	17-7
17.9	Oracle ACFS Plug-in Generic Application Programming Interface.....	17-8
17.9.1	Oracle ACFS Pre-defined Metric Types.....	17-8
17.9.2	Oracle ACFS Plug-in APIs.....	17-9
17.10	Oracle ACFS Tagging Generic Application Programming Interface.....	17-12
17.10.1	Oracle ACFS Tag Name Specifications.....	17-13
17.10.2	Oracle ACFS Tagging Error Values.....	17-13
17.10.3	acfsgettag.....	17-14
17.10.4	acfslisttags.....	17-15
17.10.5	acfsremovetag.....	17-16
17.10.6	acfssettag.....	17-17
17.11	Oracle ACFS Diagnostic Commands.....	17-18
17.11.1	acfsdbg.....	17-19
17.11.2	acfsutil blog.....	17-22
17.11.3	acfsutil dumpstate.....	17-23
17.11.4	acfsutil log.....	17-24
17.11.5	acfsutil meta.....	17-26
17.11.6	acfsutil plogconfig.....	17-28
17.12	Understanding Oracle ACFS I/O Failure Console Messages.....	17-29
18	Configuring Oracle ACFS Snapshot-Based Replication	
18.1	Configuring ssh for Use With Oracle ACFS Replication.....	18-1
18.2	Installing ssh and Cygwin on Windows.....	18-5
18.3	Upgrading to Oracle ACFS Snapshot-Based Replication.....	18-12
19	Managing Oracle Flex ASM	
19.1	Overview of Oracle Flex ASM.....	19-1
19.2	About Setting Up Oracle Flex ASM.....	19-6
19.3	Administering Oracle Flex ASM.....	19-6
20	Administering Oracle ASM Filter Driver	
20.1	About Oracle ASM Filter Driver.....	20-2
20.2	About Oracle ASMLIB and Oracle ASM Filter Driver.....	20-3

20.3	About Partition Tables and Oracle ASM Filter Driver	20-3
20.4	Configuring Oracle ASM Filter Driver	20-3
20.4.1	Configuring Oracle ASM Filter Driver During Installation	20-4
20.4.2	Configuring Oracle ASM Filter Driver After Installation	20-5
20.5	About Migrating to Oracle ASM Filter Driver From ASMLIB.....	20-7
20.5.1	Migrating Oracle ASM Disk Groups without Oracle Cluster Registry or Voting Files	20-8
20.5.2	Migrating Oracle ASM Disk Groups with Oracle Cluster Registry or Voting Files	20-8
20.5.3	Migrating Oracle ASM Disk Groups in an Oracle Grid Infrastructure Standalone (Oracle Restart) Environment	20-9
20.6	Determining Whether Oracle ASM Filter Driver Has Been Configured.....	20-10
20.7	Updating the Oracle ASM Filter Driver AFD_DISKSTRING Parameter	20-10
20.8	Updating the Oracle ASM ASM_DISKSTRING Parameter for Oracle ASM Filter Driver Disks	20-11
20.9	Setting, Clearing, and Scanning Oracle ASM Filter Driver Labels.....	20-11
20.10	Deconfiguring Oracle ASM Filter Driver	20-12
20.10.1	Deconfiguring Oracle ASM Filter Driver in an Oracle Grid Infrastructure Clusterware Environment	20-13
20.10.2	Deconfiguring Oracle ASM Filter Driver in an Oracle Grid Infrastructure Standalone (Oracle Restart) Environment	20-14
20.11	Oracle ASM Filter Driver Commands	20-15
20.11.1	afdload	20-15
20.11.2	afddriverstate	20-16

21 Creating Database Files in Oracle ASM Using SQL *Plus

21.1	Creating a Database that Uses Oracle ASM.....	21-1
21.2	Creating Tablespaces in Oracle ASM.....	21-2
21.3	Creating Redo Logs in Oracle ASM	21-2
21.4	Creating Control Files in Oracle ASM	21-3
21.4.1	Creating Control Files in Oracle ASM: Example 1	21-3
21.4.2	Creating Control Files in Oracle ASM: Example 2	21-4
21.5	Creating Archive Log Files in Oracle ASM.....	21-4

Glossary

Index

List of Examples

3-1	Sample diagnostic directory for an Oracle ASM instance.....	3-14
3-2	Determining rolling patch mode and patch level.....	3-27
4-1	Creating the DATA disk group.....	4-13
4-2	Creating the FRA disk group.....	4-13
4-3	Specifying logical and sector sizes for disks in a disk group.....	4-17
4-4	Using the QUORUM keyword when creating a failure group.....	4-19
4-5	Managing volumes with ALTER DISKGROUP VOLUME statements.....	4-21
4-6	Dropping disks from disk groups.....	4-26
4-7	Resizing disks in disk groups.....	4-28
4-8	Undropping disks in disk groups.....	4-28
4-9	Manually rebalancing a disk group.....	4-30
4-10	Querying V\$ASM_DISK for header status.....	4-34
4-11	Setting up Oracle ASM File Access Control.....	4-64
4-12	Using renamedg.....	4-72
4-13	Using CREATE DISKGROUP to create a flex disk group.....	4-74
4-14	Using ALTER DISKGROUP to migrate a normal disk group to a flex disk group.....	4-74
4-15	Adding a File Group to a Database.....	4-75
4-16	Modifying a File Group.....	4-75
4-17	Setting a file group property.....	4-78
4-18	Adding a Quota Group to a Disk Group.....	4-79
4-19	Moving a File Group to a new Quota Group.....	4-79
4-20	Using CREATE DISKGROUP with extended disk groups.....	4-81
5-1	Creating a tablespace using an alias name.....	5-7
5-2	Using an Oracle ASM File name in a SQL statement.....	5-10
5-3	Moving data files online with ALTER DATABASE MOVE DATAFILE.....	5-11
5-4	Adding an alias name for an Oracle ASM file name.....	5-14
5-5	Renaming an alias name for an Oracle ASM file name.....	5-15
5-6	Dropping an alias name for an Oracle ASM file name.....	5-15
5-7	Dropping an alias name for an Oracle ASM file name.....	5-15
5-8	Dropping a file and alias from a disk group using the alias name.....	5-15
5-9	Dropping a file and alias from a disk group using the system-generated name.....	5-15
5-10	Creating a new directory.....	5-16
5-11	Creating a new subdirectory.....	5-16
5-12	Renaming a directory.....	5-17
5-13	Dropping a directory.....	5-17
5-14	Adding a template with high redundancy to a disk group.....	5-22
5-15	Adding a template with external redundancy to a disk group.....	5-22
5-16	Modifying a disk group template.....	5-23
5-17	Dropping a template from a disk group.....	5-23
5-18	Using a user-defined template to specify attributes.....	5-23
6-1	Viewing disk group attributes with V\$ASM_ATTRIBUTE.....	6-3
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.....	6-4
6-4	Viewing disk group clients with V\$ASM_CLIENT.....	6-4
6-5	Viewing Oracle ASM File Access Control information with V\$ASM_USER.....	6-5
6-6	Viewing File Access Control information with V\$ASM_USERGROUP.....	6-5
6-7	Viewing File Access Control information with V\$ASM_USERGROUP_MEMBER.....	6-5
6-8	Viewing Oracle ASM File Access Control information with V\$ASM_FILE.....	6-5
6-9	Viewing Intelligent Data Placement information with V\$ASM_FILE.....	6-6
6-10	Viewing Intelligent Data Placement information with V\$ASM_TEMPLATE.....	6-6
8-1	Migrating the online redo logs.....	8-9

9-1	Using <code>asmca -silent -configureASM</code>	9-7
9-2	Using <code>asmca -silent -configureParameter</code>	9-8
9-3	Using <code>asmca -silent -deleteASM</code>	9-9
9-4	Using <code>asmca -silent -createDiskGroup</code>	9-10
9-5	Using <code>asmca -silent -addDisk</code>	9-11
10-1	Using wildcards with ASMCMD commands.....	10-5
10-2	Running ASMCMD commands in noninteractive mode.....	10-10
10-3	Displaying a return code when running ASMCMD in noninteractive mode.....	10-11
10-4	Running ASMCMD commands in a script.....	10-11
10-5	Displaying ASMCMD help text.....	10-11
10-6	Using the ASMCMD <code>dsget</code> command.....	10-13
10-7	Using the ASMCMD <code>dsset</code> command.....	10-14
10-8	Using the ASMCMD <code>lsct</code> command.....	10-15
10-9	Using the ASMCMD <code>lsop</code> command.....	10-15
10-10	Using the ASMCMD <code>lspwusr</code> command.....	10-16
10-11	Using the ASMCMD <code>orapwusr</code> command.....	10-17
10-12	Using the ASMCMD <code>pwcopy</code> command.....	10-18
10-13	Using the ASMCMD <code>pwcreate</code> command.....	10-19
10-14	Using the ASMCMD <code>pwdelete</code> command.....	10-19
10-15	Using the ASMCMD <code>pwget</code> command.....	10-20
10-16	Using the ASMCMD <code>pwmov</code> command.....	10-21
10-17	Using the ASMCMD <code>pwset</code> command.....	10-22
10-18	Using the ASMCMD <code>showclustermode</code> command.....	10-22
10-19	Using the ASMCMD <code>showclusterstate</code> command.....	10-22
10-20	Using the ASMCMD <code>showpatches</code> command.....	10-23
10-21	Using the ASMCMD <code>showversion</code> command.....	10-23
10-22	Using the ASMCMD <code>shutdown</code> command.....	10-24
10-23	Using the ASMCMD <code>spbackup</code> command.....	10-25
10-24	Using the ASMCMD <code>spcopy</code> command.....	10-26
10-25	Using the ASMCMD <code>spget</code> command.....	10-27
10-26	Using the ASMCMD <code>spmov</code> command.....	10-28
10-27	Using the ASMCMD <code>spset</code> command.....	10-29
10-28	Using the ASMCMD <code>startup</code> command.....	10-29
10-29	Using the ASMCMD <code>lscc</code> command.....	10-30
10-30	Using the ASMCMD <code>mkcc</code> command.....	10-32
10-31	Using the ASMCMD <code>rmcc</code> command.....	10-32
10-32	Using the <code>afd_configure</code> command.....	10-34
10-33	Using the <code>afd_deconfigure</code> command.....	10-35
10-34	Using the ASMCMD <code>afd_di</code> command.....	10-35
10-35	Using the <code>afd_dsget</code> command.....	10-36
10-36	Using the <code>afd_dsset</code> command.....	10-36
10-37	Using the <code>afd_filter</code> command.....	10-37
10-38	Using the <code>afd_label</code> command.....	10-38
10-39	Using the <code>afd_lsds</code> command.....	10-38
10-40	Using the ASMCMD <code>afd_lslbl</code> command.....	10-39
10-41	Using the ASMCMD <code>afd_refresh</code> command.....	10-40
10-42	Using the <code>afd_scan</code> command.....	10-40
10-43	Using the <code>afd_state</code> command.....	10-41
10-44	Using the <code>afd_unlabel</code> command.....	10-41
10-45	Using the ASMCMD <code>amdu_extract</code> command.....	10-43
10-46	Using the ASMCMD <code>cd</code> command.....	10-43
10-47	Using the ASMCMD <code>cp</code> command.....	10-45
10-48	Using the ASMCMD <code>du</code> command.....	10-46
10-49	Using the ASMCMD <code>find</code> command.....	10-47

10-50	Using the ASMCMD ls command.....	10-48
10-51	Using the ASMCMD lsof command.....	10-50
10-52	Using the ASMCMD mkalias command.....	10-50
10-53	Using the ASMCMD mkdir command.....	10-51
10-54	Using the ASMCMD pwd command.....	10-52
10-55	Using the ASMCMD rm command.....	10-53
10-56	Using the ASMCMD ralias command.....	10-53
10-57	Tags for the chdg XML configuration template.....	10-55
10-58	chdg sample XML configuration file	10-56
10-59	Using the ASMCMD chdg command.....	10-56
10-60	Using the ASMCMD chkdg command.....	10-57
10-61	Using the ASMCMD dropdg command.....	10-58
10-62	Using the ASMCMD iostat command.....	10-59
10-63	Using the ASMCMD lsattr command.....	10-61
10-64	Using the ASMCMD lsdg command.....	10-63
10-65	Using the ASMCMD lsdsk command.....	10-65
10-66	Using the ASMCMD lsod command.....	10-66
10-67	Using the ASMCMD md_backup command.....	10-67
10-68	Using the ASMCMD md_restore command.....	10-69
10-69	Tags for mkdg XML configuration file.....	10-70
10-70	mkdg sample XML configuration file.....	10-71
10-71	Using the ASMCMD mkdg command.....	10-71
10-72	Using the ASMCMD mount command.....	10-72
10-73	Using the ASMCMD offline command.....	10-73
10-74	Using the ASMCMD online command.....	10-74
10-75	Using the ASMCMD rebal command.....	10-74
10-76	Using the ASMCMD remap command.....	10-75
10-77	Using the ASMCMD setattr command.....	10-76
10-78	Using ASMCMD stamp.....	10-77
10-79	Using ASMCMD stamplist.....	10-78
10-80	Using the ASMCMD umount command.....	10-78
10-81	Using the chfg command.....	10-80
10-82	Using the chqg command.....	10-81
10-83	Using the lsfg command.....	10-82
10-84	Using the lsqg command.....	10-83
10-85	Using the mkfg command.....	10-84
10-86	Using the mkqg command.....	10-85
10-87	Using the mvfg command.....	10-86
10-88	Using the rmfg command.....	10-86
10-89	Using the rmqg command.....	10-87
10-90	Using the ASMCMD chtmpl command.....	10-88
10-91	Using the ASMCMD lstmpl command.....	10-89
10-92	Using the ASMCMD mktmpl command.....	10-90
10-93	Using the ASMCMD rmtmpl command.....	10-90
10-94	Using the ASMCMD chgrp command.....	10-92
10-95	Using the ASMCMD chmod command.....	10-93
10-96	Using the ASMCMD chown command.....	10-94
10-97	Using the ASMCMD groups command.....	10-95
10-98	Using the ASMCMD grpmod command.....	10-96
10-99	Using the ASMCMD lsgrp command.....	10-96
10-100	Using the ASMCMD lsusr command.....	10-97
10-101	Using the ASMCMD mkgrp command.....	10-98
10-102	Using the ASMCMD mkusr command.....	10-98
10-103	Using the ASMCMD passwd command.....	10-99

10-104	Using the ASMCMD rmgrp command.....	10-99
10-105	Using the ASMCMD rmusr command.....	10-100
10-106	Using the ASMCMD rpusr command.....	10-101
11-1	Sample audit trail file.....	11-32
12-1	Viewing encryption information in V\$ASM_ACFE_ENCRYPTION_INFO.....	12-3
12-2	Viewing security information in V\$ASM_ACFE_SECURITY_INFO.....	12-3
12-3	Viewing security rules information in V\$ASM_ACFE_SEC_RULE.....	12-3
12-4	Viewing security ruleset information in V\$ASM_ACFE_SEC_RULESET.....	12-3
12-5	Viewing security ruleset information in V\$ASM_ACFE_SEC_RULESET_RULE.....	12-4
12-6	Viewing snapshot information in V\$ASM_ACFESSNAPSHOTS.....	12-4
12-7	Viewing volume information with V\$ASM_ACFEVOLUMES.....	12-4
12-8	Viewing volume information with V\$ASM_FILESYSTEM.....	12-4
12-9	Viewing volume information with V\$ASM_VOLUME.....	12-5
12-10	Viewing volume information with V\$ASM_VOLUME_STAT.....	12-5
12-11	Viewing tag name information with V\$ASM_ACFSTAG.....	12-5
12-12	Viewing Oracle ASM information with V\$MAP_ELEMENT.....	12-7
12-13	Viewing Oracle ACFS Data File Information with V\$MAP_FILE.....	12-7
12-14	Viewing Element and File Offset Information with V\$MAP_FILE_EXTENT.....	12-7
12-15	Viewing Extent Information With V\$MAP_FILE_IO_STACK.....	12-8
14-1	Using asmca -silent -createVolume.....	14-8
14-2	Using asmca -silent -createACFS.....	14-9
15-1	Using the ASMCMD volcreate command.....	15-3
15-2	Using the ASMCMD voldelete command.....	15-5
15-3	Using the ASMCMD voldisable command.....	15-5
15-4	Using the ASMCMD volenable command.....	15-6
15-5	Using the ASMCMD volinfo command.....	15-7
15-6	Using the ASMCMD volresize command.....	15-8
15-7	Using the ASMCMD volset command.....	15-9
15-8	Using the ASMCMD volstat command.....	15-10
16-1	Displaying help for Oracle ACFS acfsutil commands.....	16-3
16-2	Using the fsck command.....	16-28
16-3	Using the mkfs command.....	16-30
16-4	Using the mkfs command to create an accelerator volume.....	16-30
16-5	Using the mount command.....	16-32
16-6	Using the umount command.....	16-33
16-7	Using the fsck command.....	16-35
16-8	Using the mkfs command.....	16-38
16-9	Using the mount command.....	16-40
16-10	Using the umount command.....	16-41
16-11	Using the fsck command.....	16-44
16-12	Using the mkfs command.....	16-46
16-13	Using the mount command.....	16-48
16-14	Using the umount command.....	16-49
16-15	Using the acfschkdsk command.....	16-51
16-16	Using the acfsdismount command.....	16-52
16-17	Using the acfsformat command.....	16-54
16-18	Using the acfsmountvol command.....	16-55
16-19	Using the acfsutil detach command.....	16-56
16-20	Using the advmutil dismount command.....	16-57
16-21	Using the advmutil list command.....	16-58
16-22	Using the advmutil mount command.....	16-59
16-23	Using the acfsutil tag info command.....	16-61
16-24	Using the acfsutil tag set command.....	16-62
16-25	Using the acfsutil tag unset command.....	16-63

16-26	Using the acfsutil repl bg command.....	16-65
16-27	Using the acfsutil repl compare command.....	16-67
16-28	Using the acfsutil repl info command.....	16-69
16-29	Using the acfsutil repl init command.....	16-72
16-30	Using the acfsutil repl pause command.....	16-73
16-31	Using the acfsutil repl resume command.....	16-73
16-32	Using the acfsutil repl sync command.....	16-75
16-33	Using the acfsutil repl terminate command.....	16-76
16-34	Using the acfsutil repl trace command.....	16-76
16-35	Using the acfsutil repl update command.....	16-78
16-36	Using the acfsutil repl upgrade command.....	16-80
16-37	Using the acfsutil audit archive command.....	16-82
16-38	Using the acfsutil audit disable command.....	16-82
16-39	Using the acfsutil audit enable command.....	16-83
16-40	Using the acfsutil audit info command.....	16-83
16-41	Using the acfsutil audit init command.....	16-84
16-42	Using the acfsutil audit purge command.....	16-85
16-43	Using the acfsutil audit read command.....	16-86
16-44	Using the acfsutil sec admin add command.....	16-88
16-45	Using the acfsutil sec admin password command.....	16-88
16-46	Using the acfsutil sec admin remove command.....	16-89
16-47	Using the acfsutil sec batch command.....	16-90
16-48	Using the acfsutil sec disable command.....	16-90
16-49	Using the acfsutil sec enable command.....	16-91
16-50	Using the acfsutil sec info command.....	16-92
16-51	Using the acfsutil sec info file command.....	16-93
16-52	Using the acfsutil sec init command.....	16-94
16-53	Using the acfsutil sec load command.....	16-96
16-54	Using the acfsutil sec prepare command.....	16-98
16-55	Using the acfsutil sec realm add command.....	16-101
16-56	Using the acfsutil sec realm audit disable command.....	16-102
16-57	Auditing Oracle ACFS security backup operators.....	16-104
16-58	Auditing only security realm users.....	16-104
16-59	Running acfsutil sec realm audit enable multiple times.....	16-104
16-60	Running acfsutil sec realm audit info.....	16-105
16-61	Using the acfsutil sec realm clone command.....	16-106
16-62	Using the acfsutil sec realm create command.....	16-108
16-63	Using the acfsutil sec realm delete command.....	16-109
16-64	Using the acfsutil sec realm destroy command.....	16-109
16-65	Using the acfsutil sec rule clone command.....	16-110
16-66	Using the acfsutil sec rule create command.....	16-111
16-67	Using the acfsutil sec rule destroy command.....	16-112
16-68	Using the acfsutil sec rule edit command.....	16-113
16-69	Using the acfsutil sec ruleset clone command.....	16-114
16-70	Using the acfsutil sec ruleset create command.....	16-115
16-71	Using the acfsutil sec ruleset destroy command.....	16-116
16-72	Using the acfsutil sec ruleset edit command.....	16-116
16-73	Using the acfsutil sec save command.....	16-117
16-74	Using the acfsutil encr info command.....	16-119
16-75	Using the acfsutil encr init command.....	16-120
16-76	Using the acfsutil encr off command.....	16-120
16-77	Using the acfsutil encr on command.....	16-122
16-78	Using the acfsutil encr rekey command.....	16-123
16-79	Using the acfsutil encr set command.....	16-125

16-80	Using the acfsutil keystore migrate command.....	16-125
16-81	Using the acfsutil snap convert command.....	16-127
16-82	Using the acfsutil snap create command on Linux.....	16-129
16-83	Using the acfsutil snap create command on Windows.....	16-129
16-84	Using the acfsutil snap delete command.....	16-131
16-85	Using the acfsutil snap duplicate apply command.....	16-132
16-86	Using the acfsutil snap duplicate create command.....	16-134
16-87	Using the acfsutil snap info command on Linux.....	16-135
16-88	Using the acfsutil snap info command on Windows.....	16-136
16-89	Using the acfsutil snap quota command.....	16-138
16-90	Using the acfsutil snap remaster command.....	16-139
16-91	Using the acfsutil snap rename command.....	16-140
16-92	Using the acfsutil compress copy command.....	16-141
16-93	Using the acfsutil compress info command.....	16-141
16-94	Using the acfsutil compress off command.....	16-142
16-95	Using the acfsutil compress on command.....	16-143
16-96	Using the acfsutil accel replace command.....	16-144
16-97	Using the acfsutil compat get command.....	16-145
16-98	Using the acfsutil compat set command.....	16-146
16-99	Using the acfsutil defrag dir command.....	16-146
16-100	Using the acfsutil defrag file command.....	16-147
16-101	Using the acfsutil info file command.....	16-148
16-102	Using the acfsutil info fs command on Linux.....	16-152
16-103	Using the acfsutil info fs command on Windows.....	16-153
16-104	Using the acfsutil info fs command with the -s option.....	16-154
16-105	Using the acfsutil info id command.....	16-155
16-106	Disabling the Oracle ACFS driver.....	16-155
16-107	Enabling the Oracle ACFS driver for storage visibility: poll model.....	16-157
16-108	Enabling the Oracle ACFS driver for storage visibility: post on interval model.....	16-158
16-109	Displaying information about the Oracle ACFS driver.....	16-159
16-110	Using the acfsutil registry command.....	16-162
16-111	Using the acfsutil rmfs command.....	16-163
16-112	Using the Oracle ACFS acfsutil scrub command.....	16-164
16-113	Using the acfsutil size command.....	16-167
16-114	Using the acfsutil tune command.....	16-168
16-115	Using advmutil canonical.....	16-169
16-116	Using advmutil tune.....	16-170
16-117	Using advmutil volinfo.....	16-171
17-1	Application Plug-in for Storage Visibility: Poll Model.....	17-11
17-2	Application Plug-in for File Content: Post Model.....	17-11
17-3	Application for Resolving the File Path from a Fenum and Genum Pair.....	17-12
17-4	Retrieving a file tag value.....	17-15
17-5	Listing file tags.....	17-16
17-6	Removing file tags.....	17-17
17-7	Setting file tags.....	17-18
17-8	Using the acfsdbg command.....	17-22
17-9	Using acfsutil blog.....	17-23
17-10	Using the acfsutil dumpstate command.....	17-24
17-11	Using acfsutil log.....	17-26
17-12	Using the acfsutil meta command.....	17-28
17-13	Using the Oracle ACFS acfsutil plogconfig command.....	17-29
21-1	Creating control files in Oracle ASM using NORESETLOGS.....	21-3
21-2	Creating control files in Oracle ASM using RESETLOGS.....	21-4
21-3	Archived log files in an Oracle ASM disk group.....	21-5

List of Figures

1-1	Oracle ASM for Single-Instance Oracle Databases	1-4
1-2	Oracle ASM Cluster Configuration with Oracle RAC.....	1-5
1-3	Oracle ASM Cluster with Single-Instance Oracle Databases.....	1-5
1-4	Oracle ASM File Allocation in a Disk Group.....	1-9
1-5	Oracle ASM Fine-Grained Striping.....	1-10
1-6	Oracle ASM Coarse-Grained Striping.....	1-11
4-1	Oracle ASM File Groups.....	4-75
5-1	Hierarchical Structure of Oracle ASM Folders in an XML DB Installation.....	5-18
11-1	Oracle ACFS Storage Layers.....	11-2
19-1	Oracle Flex ASM Client Configurations.....	19-2
19-2	Oracle ACFS and Oracle ADVM in Oracle Flex ASM Configuration.....	19-4
19-3	Oracle IOServer Configuration.....	19-5

List of Tables

3-1	One operating system group and one set of privileges for all Oracle ASM users.....	3-29
3-2	Separated operating system groups and privileges for Oracle ASM users.....	3-30
4-1	Possible combinations of sector size and logical sector size by disk size and type.....	4-16
4-2	Mirroring options for Oracle ASM disk group types.....	4-38
4-3	Examples of disk group compatibility attribute settings.....	4-53
4-4	Oracle ASM features enabled by disk group compatibility attribute settings.....	4-58
4-5	Oracle ACFS features enabled by disk group compatibility attribute settings.....	4-59
4-6	Maximum Oracle ASM file sizes for disk groups with AU_SIZE set to 1 MB.....	4-60
4-7	Maximum Oracle ASM file sizes for disk groups with AU_SIZE set to 4 MB.....	4-60
5-1	File types supported by Oracle ASM.....	5-1
5-2	Valid usage for Oracle ASM file name forms.....	5-4
5-3	Oracle file types and Oracle ASM file type tags.....	5-5
5-4	Permitted values for Oracle ASM template striping attribute.....	5-20
5-5	Permitted values for Oracle ASM template redundancy attribute.....	5-20
5-6	Oracle ASM system default templates attribute settings.....	5-21
5-7	Permitted values for Oracle ASM template primary and mirror region attribute.....	5-22
6-1	Oracle ASM dynamic views for disk group information.....	6-1
9-1	Options for the ASMCA command-line tool.....	9-4
9-2	Options for the -configureASM command.....	9-6
9-3	Options for the -configureParameter command.....	9-8
9-4	Options for the -deleteASM command.....	9-8
9-5	Options for the -createDiskGroup command.....	9-9
9-6	Options for the -addDisk command.....	9-11
9-7	Options for the -editDiskGroupAttributes command.....	9-11
9-8	Options for the —createLabel command.....	9-12
9-9	Options for the —dropLabel command.....	9-12
9-10	Options for the -createFileGroups command.....	9-13
9-11	Options for the -dropFileGroups command.....	9-13
9-12	Options for the -createQuotaGroups command.....	9-14
9-13	Options for the -dropQuotaGroups command.....	9-14
10-1	Types of ASMCMD commands.....	10-2
10-2	ASMCMD interactive options.....	10-8
10-3	ASMCMD return codes.....	10-10
10-4	Summary of ASMCMD instance management commands.....	10-12
10-5	Options for the dsget command.....	10-13
10-6	Options for the dsset command.....	10-14
10-7	Options for the lsct command.....	10-15
10-8	Options for the lspwusr command	10-16
10-9	Options for the orapwusr command	10-16
10-10	Options for the pwcopu command.....	10-17
10-11	Options for the pwcreate command.....	10-18
10-12	Options for the pwdelete command.....	10-19
10-13	Options for the pwget command.....	10-20
10-14	Options for the pwmove command.....	10-20
10-15	Options for the pwset command.....	10-21
10-16	Options for the showversion command.....	10-23
10-17	Options for the shutdown command.....	10-24
10-18	Options for the spbackup command.....	10-25
10-19	Options for the spcopy command.....	10-26
10-20	Options for the spmove command.....	10-28
10-21	Options for the spset command.....	10-28

10-22	Options for the startup command.....	10-29
10-23	Summary of ASMCMD client cluster management commands.....	10-30
10-24	Options for the lsccl command	10-30
10-25	Options for the mkcc command	10-31
10-26	Options for the rmcc command	10-32
10-27	Summary of ASMCMD Oracle ASM Filter Driver management commands.....	10-33
10-28	Options for the afd_configure command.....	10-33
10-29	Options for the afd_deconfigure command.....	10-34
10-30	Options for the afd_di command.....	10-35
10-31	Options for the afd_dsset command.....	10-36
10-32	Options for the afd_filter command.....	10-37
10-33	Options for the afd_label command.....	10-37
10-34	Options for the afd_lsdsk command.....	10-38
10-35	Options for the afd_lslbl command.....	10-39
10-36	Options for the afd_refresh command.....	10-39
10-37	Options for the afd_scan command.....	10-40
10-38	Options for the afd_state command.....	10-41
10-39	Options for the afd_unlabel command.....	10-41
10-40	Summary of ASMCMD file management commands.....	10-42
10-41	Options for the amdu_extract command	10-42
10-42	Options for the cd command	10-43
10-43	Options for the cp command	10-44
10-44	Options for the du command.....	10-46
10-45	Options for the find command	10-46
10-46	Options for the ls command.....	10-47
10-47	Options for the lsof command.....	10-49
10-48	Options for the mkalias command.....	10-50
10-49	Options for the mkdir command.....	10-51
10-50	Options for the rm command.....	10-52
10-51	Options for the rmalias command.....	10-53
10-52	Summary of ASMCMD Disk group management commands.....	10-54
10-53	Options for the chdg command	10-55
10-54	Options for the chkdg command	10-57
10-55	Options for the dropdg command	10-57
10-56	Options for the iostat command.....	10-58
10-57	Attribute descriptions for iostat command output.....	10-58
10-58	Options for the lsattr command.....	10-60
10-59	Options for the lsdg command.....	10-61
10-60	Attribute descriptions for lsdg command output.....	10-62
10-61	Options for the lsdsk command.....	10-63
10-62	Options for the lsod command.....	10-65
10-63	Options for the md_backup command.....	10-66
10-64	Options for the md_restore command.....	10-68
10-65	Options for the mkdg command	10-70
10-66	Options for the mount command	10-72
10-67	Options for the offline command.....	10-72
10-68	Options for the online command.....	10-73
10-69	Options for the rebal command.....	10-74
10-70	Options for the remap command.....	10-75
10-71	Options for the setattr command.....	10-76
10-72	Options for the stamp command.....	10-76
10-73	Options for the stampist command.....	10-77
10-74	Options for the umount command	10-78
10-75	Summary of ASMCMD file group commands.....	10-79

10-76	Options for the chfg command.....	10-80
10-77	Options for the chqg command.....	10-81
10-78	Options for the lsfg command.....	10-81
10-79	Options for the lsqg command.....	10-82
10-80	Options for the mkfg command.....	10-83
10-81	Options for the mkqg command.....	10-85
10-82	Options for the mvfg command.....	10-85
10-83	Options for the rmfg command.....	10-86
10-84	Options for the rmqg command.....	10-87
10-85	Summary of ASMCMD template management commands.....	10-87
10-86	Options for the chtmpl command	10-87
10-87	Options for the lstmpl command.....	10-88
10-88	Options for the mktmpl command	10-89
10-89	Options for the rmtmpl command	10-90
10-90	Summary of ASMCMD file access control commands.....	10-91
10-91	Options for the chgrp command	10-92
10-92	Options for the chmod command	10-92
10-93	Options for the chown command	10-94
10-94	Options for the groups command	10-94
10-95	Options for the grpmod command	10-95
10-96	Options for the lsgrp command.....	10-96
10-97	Options for the lsusr command.....	10-97
10-98	Options for the mkgrp command	10-97
10-99	Options for the mkusr command.....	10-98
10-100	Options for the passwd command.....	10-98
10-101	Options for the rmgrp command	10-99
10-102	Options for the rmusr command.....	10-100
10-103	Options for the rpusr command.....	10-100
12-1	Oracle ASM dynamic views for Oracle ACFS information.....	12-1
14-1	Options for the -createACFSSnapshot command.....	14-7
14-2	Options for the -deleteACFSSnapshot command.....	14-7
14-3	Options for the -createVolume command.....	14-8
14-4	Options for the -createACFS command.....	14-9
15-1	Summary of ASMCMD volume management commands.....	15-1
15-2	Options for the volcreate command.....	15-2
15-3	Options for the voldelete command.....	15-4
15-4	Options for the voldisable command.....	15-5
15-5	Options for the volenable command.....	15-6
15-6	Options for the volinfo command.....	15-6
15-7	Options for the volresize command.....	15-8
15-8	Options for the volset command.....	15-8
15-9	Options for the volstat command.....	15-9
16-1	Summary of Oracle ACFS commands for Linux environments.....	16-26
16-2	Options for the Linux fsck command.....	16-26
16-3	Options for the Linux mkfs command.....	16-28
16-4	Options for the Linux mount command.....	16-31
16-5	Options for the Linux umount command.....	16-32
16-6	Summary of Oracle ACFS commands for Solaris.....	16-33
16-7	Options for the Solaris fsck command.....	16-34
16-8	Options for the Solaris mkfs command.....	16-36
16-9	Options for the Solaris mount command.....	16-38
16-10	Options for the Solaris umount command.....	16-41
16-11	Summary of Oracle ACFS commands for AIX.....	16-42
16-12	Options for the AIX fsck command.....	16-43

16-13	Options for the AIX mkfs command.....	16-45
16-14	Options for the AIX mount command.....	16-47
16-15	Options for the AIX umount command.....	16-48
16-16	Summary of Oracle ACFS file system commands for Windows.....	16-49
16-17	Summary of Non-Oracle ACFS file system commands for Windows.....	16-49
16-18	Options for the acfschkdsk command.....	16-50
16-19	Options for the acfsdismount command.....	16-52
16-20	Options for the acfsformat command.....	16-53
16-21	Options for the acfsmountvol command.....	16-55
16-22	Options for the advmutil dismount command.....	16-57
16-23	Options for the advmutil list command.....	16-58
16-24	Options for the advmutil mount command.....	16-59
16-25	Summary of commands for Oracle ACFS tagging.....	16-60
16-26	Options for the acfsutil tag info command.....	16-61
16-27	Options for the acfsutil tag set command.....	16-62
16-28	Options for the acfsutil tag unset command.....	16-63
16-29	Summary of commands for Oracle ACFS replication.....	16-64
16-30	Options for the acfsutil repl bg command.....	16-64
16-31	Options for the acfsutil repl compare command.....	16-66
16-32	Options for the acfsutil repl info command.....	16-67
16-33	Options for the acfsutil repl init command.....	16-70
16-34	Options for the acfsutil repl pause command.....	16-72
16-35	Options for the acfsutil repl resume command.....	16-73
16-36	Options for the acfsutil repl sync command.....	16-74
16-37	Options for the acfsutil repl terminate command.....	16-75
16-38	Options for the acfsutil repl trace command.....	16-76
16-39	Options for the acfsutil repl update command.....	16-77
16-40	Options for the acfsutil repl upgrade command.....	16-79
16-41	Summary of commands for Oracle ACFS auditing.....	16-81
16-42	Options for the acfsutil audit archive command.....	16-81
16-43	Options for the acfsutil audit disable command.....	16-82
16-44	Options for the acfsutil audit enable command.....	16-83
16-45	Options for the acfsutil audit info command.....	16-83
16-46	Options for the acfsutil audit init command.....	16-84
16-47	Options for the acfsutil audit purge command.....	16-85
16-48	Options for the acfsutil audit read command.....	16-85
16-49	Summary of commands for Oracle ACFS security.....	16-86
16-50	Options for the acfsutil sec admin add command.....	16-87
16-51	Options for the acfsutil sec admin remove command.....	16-89
16-52	Options for the acfsutil sec batch command.....	16-89
16-53	Options for the acfsutil sec disable command.....	16-90
16-54	Options for the acfsutil sec enable command.....	16-91
16-55	Options for the acfsutil sec info command.....	16-92
16-56	Options for the acfsutil sec info file command.....	16-93
16-57	Options for the acfsutil sec init command.....	16-93
16-58	Options for the acfsutil sec load command.....	16-95
16-59	Options for the acfsutil sec prepare command.....	16-96
16-60	Options for the acfsutil sec realm add command.....	16-98
16-61	Security Realm Command Rules.....	16-100
16-62	Options for the acfsutil sec realm audit disable command.....	16-102
16-63	Options for the acfsutil sec realm audit enable command.....	16-103
16-64	Options for the acfsutil sec realm audit info command.....	16-105
16-65	Options for the acfsutil sec realm clone command.....	16-106
16-66	Options for the acfsutil sec realm create command.....	16-107

16-67	Options for the acfsutil sec realm delete command.....	16-108
16-68	Options for the acfsutil sec realm destroy command.....	16-109
16-69	Options for the acfsutil sec rule clone command.....	16-110
16-70	Options for the acfsutil sec rule create command.....	16-111
16-71	Options for the acfsutil sec rule destroy command.....	16-112
16-72	Options for the acfsutil sec rule edit command.....	16-113
16-73	Options for the acfsutil sec ruleset clone command.....	16-113
16-74	Options for the acfsutil sec ruleset create command.....	16-114
16-75	Options for the acfsutil sec ruleset destroy command.....	16-115
16-76	Options for the acfsutil sec ruleset edit command.....	16-116
16-77	Options for the acfsutil sec save command.....	16-117
16-78	Summary of commands for Oracle ACFS encryption.....	16-117
16-79	Options for the acfsutil encr info command.....	16-118
16-80	Options for the acfsutil encr init command.....	16-119
16-81	Options for the acfsutil encr off command.....	16-120
16-82	Options for the acfsutil encr on command.....	16-121
16-83	Options for the acfsutil encr rekey command.....	16-122
16-84	Options for the acfsutil encr set command.....	16-123
16-85	Options for the acfsutil keystore migrate command.....	16-125
16-86	Summary of commands for Oracle ACFS snapshots.....	16-125
16-87	Options for the acfsutil snap convert command.....	16-126
16-88	Options for the acfsutil snap create command.....	16-127
16-89	Options for the acfsutil snap delete command.....	16-130
16-90	Options for the acfsutil snap duplicate apply command.....	16-131
16-91	Options for the acfsutil snap duplicate create command.....	16-133
16-92	Options for the acfsutil snap info command.....	16-134
16-93	Options for the acfsutil snap quota command.....	16-137
16-94	Options for the acfsutil snap remaster command.....	16-138
16-95	Options for the acfsutil snap rename command.....	16-139
16-96	Summary of commands for Oracle ACFS compression.....	16-140
16-97	Options for the acfsutil compress copy command.....	16-140
16-98	Options for the acfsutil compress info command.....	16-141
16-99	Options for the acfsutil compress off command.....	16-142
16-100	Options for the acfsutil compress on command.....	16-142
16-101	Summary of Oracle ACFS command-line utilities.....	16-143
16-102	Options for the acfsutil accel replace command.....	16-144
16-103	Options for the acfsutil compat get command.....	16-145
16-104	Options for the acfsutil compat set command.....	16-145
16-105	Options for the acfsutil defrag dir command.....	16-146
16-106	Options for the acfsutil defrag file command.....	16-147
16-107	Options for the acfsutil info file command.....	16-147
16-108	Options for the acfsutil info fs command.....	16-149
16-109	Options for the acfsutil info id command.....	16-154
16-110	Options for the acfsutil plugin disable command.....	16-155
16-111	Options for the acfsutil plugin enable command.....	16-156
16-112	Options for the acfsutil plugin info command.....	16-159
16-113	Options for the acfsutil registry command.....	16-160
16-114	Options for the acfsutil rmfs command.....	16-162
16-115	Options for the acfsutil scrub command.....	16-163
16-116	Options for the acfsutil size command.....	16-165
16-117	Options for the acfsutil tune command.....	16-168
16-118	Options for the advmutil canonical command.....	16-169
16-119	Options for the advmutil tune command.....	16-170
16-120	Options for the advmutil volinfo command.....	16-170

17-1	Maximum file sizes for Oracle ADVM volumes.....	17-2
17-2	Options for the acfsload command.....	17-7
17-3	Options for the acfsdriverstate command.....	17-7
17-4	Options for the acfsgettag command.....	17-14
17-5	Options for the acfslisttags command.....	17-15
17-6	Options for the acfsremovetag command.....	17-17
17-7	Options for the acfssettag command.....	17-17
17-8	Summary of Oracle ACFS diagnostic commands.....	17-19
17-9	Options for the acfsdbg command.....	17-19
17-10	Subcommands for acfsdbg.....	17-20
17-11	Options for the acfsutil blog command.....	17-23
17-12	Options for the acfsutil dumpstate command.....	17-24
17-13	Options for the acfsutil log command.....	17-24
17-14	Options for the acfsutil meta command.....	17-26
17-15	Options for the acfsutil plogconfig command.....	17-28
20-1	Options for the afdload command.....	20-15
20-2	Options for the afddriverstate command.....	20-16

Preface

The *Oracle Automatic Storage Management Administrator's Guide* describes how to administer Oracle Automatic Storage Management (Oracle ASM) for Oracle Database.

This Preface contains the following topics:

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

Audience

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

- Administer and manage Oracle ASM
- Configure and administer Oracle 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](#) (page xxxiv)".

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, refer to the following Oracle resources:

- *Oracle Database Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle Database Reference*
- *Oracle Database SQL Language Reference*
- *Oracle Database 2 Day + Real Application Clusters Guide*
- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Real Application Clusters Administration and Deployment Guide*
- Platform-specific guides, including Oracle Database, Oracle Grid Infrastructure, and Oracle Real Application Clusters installation guides

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.

Changes in This Release for Oracle Automatic Storage Management Administrator's Guide

This chapter describes changes in *Oracle Automatic Storage Management Administrator's Guide*.

- [Changes in Oracle Automatic Storage Management 12c Release 2 \(12.2.0.1\)](#) (page xxxv)
- [Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 2 \(12.2.0.1\)](#) (page xxxvii)
- [Changes in Oracle Automatic Storage Management 12c Release 1 \(12.1.0.2\)](#) (page xl)
- [Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 1 \(12.1.0.2\)](#) (page xli)
- [Changes in Oracle Automatic Storage Management 12c Release 1 \(12.1.0.1\)](#) (page xlii)
- [Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 1 \(12.1.0.1\)](#) (page xlvi)

See Also:

Oracle Database New Features Guide for a complete description of the new features in Oracle Database 12c Release 1 (12.1) and Oracle Database 12c Release 2 (12.2).

Changes in Oracle Automatic Storage Management 12c Release 2 (12.2.0.1)

The following are changes in Oracle Automatic Storage Management 12c Release 2 (12.2.0.1).

- [New Features](#) (page xxxvi)
- [Deprecated Features](#) (page xxxvii)

New Features

These are new features for Oracle ASM 12c Release 2 (12.2.0.1):

- Oracle ASM Flex Disk Groups
Oracle ASM provides database-oriented storage management with Oracle ASM flex and extended disk groups. For information about Oracle ASM flex disk groups, extended disk groups, and file groups, refer to [Managing Oracle ASM Flex Disk Groups](#) (page 4-72).
- Prioritized Rebalancing for Oracle ASM File Groups
Oracle ASM file groups have the ability to set priority levels for rebalancing with the `PRIORITY` property.
For information about Oracle ASM file group properties, refer to [About Oracle ASM File Group Properties](#) (page 4-76).
- Oracle ASM Support for Preferred Read on Extended Clusters
An Oracle extended cluster is special purpose Oracle flex cluster which contains nodes that span multiple physically separated locations. An Oracle flex cluster is a scalable cluster in which servers can be assigned specific roles for database or application functions.
In an Oracle extended cluster, the `PREFERRED_READ.ENABLED` disk group attribute controls whether preferred read functionality is enabled for a disk group. If preferred read functionality is enabled, then this functionality enables an instance to determine and read from disks at the same site as itself, which can improve performance. For information about the `PREFERRED_READ.ENABLED` disk group attribute, refer to [PREFERRED_READ.ENABLED](#) (page 4-7).
- Oracle ASM Extended Support for 4K Sector Size
This features provides the new disk group attribute `LOGICAL_SECTOR_SIZE`. This attribute sets the logical sector size of the disk group instead of the size that is estimated from the disks that are added to the disk group. The value specifies the smallest possible I/O that the disk group can accept, and also limits what kind of disks can join the disk group. A new column `LOGICAL_SECTOR_SIZE` is included in the `V$ASM_DISKGROUP` and `V$ASM_DISK` views. This column represents the logical sector size value of the disk group or disk in bytes.
For information about extended support for 4K sector size, refer to [Specifying the Logical and Physical Sector Size](#) (page 4-15).
- Oracle IO Server
An Oracle IO Server (IOS) instance provides Oracle ASM file access for Oracle Database instances on nodes of Oracle member clusters that do not have connectivity to Oracle ASM managed disks.
For information about remote Oracle IO Server and Oracle Flex ASM, refer to [Managing Oracle Flex ASM](#) (page 19-1).
- Updates for Oracle ASM Filter Driver Installation and Configuration
The installation and configuration for Oracle ASM Filter Driver (Oracle ASMFDF) can now be enabled as an automated process during Oracle Grid Infrastructure installation.

Oracle ASMFD is available on Solaris systems starting with Oracle Database 12c Release 2 (12.2.0.1).

For more information about Oracle ASMFD, refer to "[Administering Oracle ASM Filter Driver](#) (page 20-1)".

Deprecated Features

These are deprecated features for Oracle ASM 12c Release 2 (12.2.0.1).

- Deprecation of Oracle ASM Intelligent Data Placement

The Oracle ASM Intelligent Data Placement (IDP) feature has been deprecated in Oracle ASM 12c Release 2 (12.2) and may be desupported in a future release.

For information about IDP, refer to [Intelligent Data Placement](#) (page 4-26).

- Deprecation of ASM_PREFERRED_READ_FAILURE_GROUPS Initialization Parameter

The ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter has been deprecated in Oracle ASM 12c Release 2 (12.2) and may be desupported in a future release.

In Oracle ASM 12c Release 2 (12.2), specifying the preferred read failure groups is done automatically, so the use of the ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter is no longer required. Starting with Oracle ASM 12c Release 2, the PREFERRED_READ.ENABLED disk group attribute controls whether the preferred read functionality is enabled.

For information about the ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter, refer to [ASM_PREFERRED_READ_FAILURE_GROUPS](#) (page 3-12). For information about the PREFERRED_READ.ENABLED disk group attribute, refer to [PREFERRED_READ.ENABLED](#) (page 4-7).

Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 2 (12.2.0.1)

The following are changes to Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) for Oracle ASM 12c Release 2 (12.2.0.1).

New Features

These are new features for Oracle ACFS and Oracle ADVM in Oracle ASM 12c Release 2 (12.2.0.1).

- Oracle ACFS Snapshot-Based Replication

Oracle ACFS snapshot-based replication transfers the differences between successive snapshots of the primary file system to the standby file system using the standard ssh protocol.

For an overview of Oracle ACFS snapshot-based replication, refer to [Oracle ACFS Replication](#) (page 11-36). For information about Oracle ACFS replication commands, refer to [Oracle ACFS Command-Line Tools for Replication](#) (page 16-63).

For information about converting an existing Oracle ACFS replication environment to snapshot-based replication, refer to [Configuring Oracle ACFS Snapshot-Based Replication](#) (page 18-1).

Note:

For information about the previous Oracle ACFS replication system available in Oracle Grid Infrastructure release 1 (12.1), refer to *Oracle Automatic Storage Management Administrator's Guide* in the [Database Online Documentation 12c Release 1 \(12.1\)](#) library.

- Oracle ACFS Snapshot Enhancements

Oracle ACFS snapshot enhancements include `acfsutil` operating system commands to create and manage snapshot duplicates. Additional commands provide support for remastering, renaming, and setting quotas for snapshots.

For information about Oracle ACFS snapshots, refer to [About Oracle ACFS Snapshots](#) (page 11-11). For information about `acfsutil` commands to manage Oracle ACFS snapshots, refer to [Oracle ACFS Command-Line Tools for Snapshots](#) (page 16-125).

- Oracle ACFS Compression

Oracle ACFS provides file system compression functionality with `acfsutil compress` commands.

For information about Oracle ACFS compression, refer to [Oracle ACFS Compression](#) (page 11-28).

- Oracle ACFS Defragger

Databases that share storage with snapshots or with the base of the file system can become fragmented under active online transaction processing (OLTP) workloads. This fragmentation can cause the location of the data in the volume to be discontinuous for sequential scans. Oracle ACFS automatically defragments these files in the background. Fragmentation is reported through the `acfsutil defrag dir` and `file` commands, and it can also be viewed with the `acfsutil info file` command. The `acfsutil defrag dir` and `file` commands also enable on-demand defragmentation of a file in the event the automatic defragmentation does not occur quickly enough. Defragmentation may also be needed when database data files reside in file systems enabled with Oracle ACFS compression.

For more information about database files on Oracle ACFS, refer to "[About Oracle ACFS and Database Data Files](#) (page 11-6)". For information about the `acfsutil defrag` and `acfsutil info` commands, refer to [Oracle ACFS Command-Line Utilities](#) (page 16-143).

- Oracle ACFS Support for 4K Sectors

Oracle ACFS supports I/O requests in multiples of 4K logical sector sizes as well as continued support for 512-byte logical sector size I/O requests. The `i 4096` option is provided with the `acfsformat` command on Windows and the `mkfs` command in AIX, Linux and Solaris environments.

For information about the Oracle ACFS `mkfs` command on Linux, refer to [mkfs](#) (page 16-28). For information about the Oracle ACFS `acfsformat` command, refer to [acfsformat](#) (page 16-53).

- Oracle ACFS Automatic Resize

Oracle ACFS provides an automatic resize option with the `acfsutil size` command that enables you to specify an increment by which an Oracle ACFS file system grows automatically if the amount of available free space in the file system falls below a specified amount. There is also an option to specify the maximum size allowed when using the automatic resize option. The output of the `acfsutil info fs` command displays the automatic resize increment and maximum amounts.

For more information about the `acfsutil size` command, refer to [acfsutil size](#) (page 16-164).

- Oracle ACFS Metadata Acceleration

Oracle ACFS supports accelerator metadata storage. This support enables many critical Oracle ACFS metadata structures, including extent metadata, storage bitmaps, volume logs, and some snapshot metadata to be placed on accelerator storage.

The accelerator volume can be created on Linux environments with the `-a` option of the `mkfs` command. To use the `-a` option, the value of `COMPATIBLE.ADVM` must be at least `12.2`.

For more information about the metadata storage, refer to [Oracle ACFS Accelerator Volume](#) (page 11-43). For information about specifying accelerator storage with the `mkfs` command, refer to [mkfs](#) (page 16-28). For commands that can be used to view information about accelerator storage, refer to [acfsutil info file](#) (page 16-147) and [acfsutil info fs](#) (page 16-149).

- Oracle ACFS NAS Maximum Availability eXtensions

Oracle ACFS provides enhancements to High Availability Common Internet File System (HACIFS) with Oracle ACFS NAS Maximum Availability eXtensions (Oracle ACFS NAS MAX).

For more information, refer to [Oracle ACFS NAS Maximum Availability eXtensions](#) (page 11-44).

- Oracle ACFS Plugins for File Content Data Collection

Oracle ACFS plugins support file content data collection. Both polling and interval based capture are supported with the file content data collection.

For more information about Oracle ACFS plug-ins and the `acfsfileid_lookup` API, refer to [Oracle ACFS Plug-in Generic Application Programming Interface](#) (page 17-8).

- Oracle ACFS Sparse Files

Oracle ACFS provides support for sparse files. Oracle ACFS sparse files greatly benefit NFS client write operations which are commonly received out of order by the NFS server and the associated Oracle ACFS file system.

Usually when an application writes beyond the end of file, storage is allocated and zeroes inserted for the region of the file beyond the old end of file and the beginning of the new write. With this feature, rather than inserting zeroes, a hole remains in the file, which benefits NFS performance and also the performance and disk utilization of other applications that perform this type of writing intentionally. This feature also has storage savings benefits for files that are inherently sparse, meaning they have a lot of unused space, such as some image

files for virtual machines. Oracle ACFS fills the holes in the file with zeroes in memory when the holes are read, rather than allocating storage for the zeroes.

To use this feature, the `COMPATIBLE .ADVM` disk group attribute must be set to `12.2` or greater.

For an overview of Oracle ACFS storage management, refer to [Overview of Oracle ACFS](#) (page 11-1).

- Oracle ACFS Scrubbing Functionality

Oracle ACFS provides scrubbing functionality with the `acfsutil scrub` command to check for and report any inconsistencies in the metadata or file data.

For information about the `acfsutil scrub` command, refer to [Oracle ACFS Command-Line Utilities](#) (page 16-143).

- Oracle ACFS Loopback Functionality

Oracle ACFS supports loopback functionality on the Linux operating system, enabling Oracle ACFS files to be accessed as devices.

For information about Oracle ACFS loopback support, refer to [Oracle ACFS Loopback Support](#) (page 17-4).

- Oracle ACFS Diagnostic Commands

Oracle ACFS provides various `acfsutil` commands for diagnostic purposes. These commands are intended to be run only when requested by Oracle Support to generate diagnostic data for analysis.

For information about the diagnostic commands, refer to [Oracle ACFS Diagnostic Commands](#) (page 17-18).

Changes in Oracle Automatic Storage Management 12c Release 1 (12.1.0.2)

The following are changes in Oracle Automatic Storage Management 12c Release 1 (12.1.0.2).

New Features

These are new features for Oracle ASM 12c Release 1 (12.1.0.2).

- Oracle ASM Filter Driver

Oracle ASM Filter Driver (Oracle ASMFD) is a kernel module that resides in the I/O path of the Oracle ASM disks. Oracle ASM uses the filter driver to validate write I/O requests to Oracle ASM disks.

The Oracle ASM Filter Driver rejects any I/O requests that are invalid. This action eliminates accidental overwrites of Oracle ASM disks that would cause corruption in the disks and files within the disk group. For example, the Oracle ASM Filter Driver filters out all non-Oracle I/Os which could cause accidental overwrites.

For more information about Oracle ASMFD, refer to [Administering Oracle ASM Filter Driver](#) (page 20-1).

Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 1 (12.1.0.2)

The following are changes for Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) in Oracle ASM 12c Release 1 (12.1.0.2).

New Features

These are new features for Oracle ACFS and Oracle ADVM in Oracle ASM 12c Release 1 (12.1.0.2).

- Oracle ACFS Support For Database Data Files on Windows

Oracle ACFS in Oracle Grid 12c Release 1 (12.1.0.2) on Windows supports all database files, except for data files and redo logs in an Oracle Restart (standalone server) configuration. For support of database files on Windows, the `COMPATIBLE .ADVM` attribute must be set to 12.1.0.2.

For more information about database files on Oracle ACFS, refer to [About Oracle ACFS and Database Data Files](#) (page 11-6).

- Oracle Exadata Support for Oracle ACFS and Oracle ACFS Support for General Purpose Files

Oracle Exadata on Linux supports Oracle ACFS in Oracle Grid 12c Release 1 (12.1.0.2). Oracle ACFS supports general purpose files as well as database files; however, Oracle ASM is the best practice for support of database files.

Oracle ACFS can be configured for use with the database, particularly to leverage Oracle ACFS snapshots for database testing and development. However, Oracle ACFS does not currently have the ability to push database operations directly into storage.

For more information about database files on Oracle ACFS, refer to [About Oracle ACFS and Database Data Files](#) (page 11-6).

- Oracle ACFS Support for 1023 Snapshots

A total of 1023 read-only, read-write, or combination of read-only and read-write snapshot views are supported for each file system. For information about Oracle ACFS snapshots, refer to [About Oracle ACFS Snapshots](#) (page 11-11).

- Oracle ACFS Support for Oracle Database File Mapping Views

Oracle ACFS supports Oracle Database file mapping views to the Oracle ASM device level. For information about file mapping views related to Oracle ACFS, refer to [Oracle ACFS Support for Oracle Database File Mapping Views](#) (page 12-6).

- Oracle ACFS Encryption Support for Oracle Key Vault

Oracle ACFS encryption supports Oracle Key Vault as a key store. For information about using this feature with Oracle ACFS encryption, refer to [Oracle ACFS Encryption](#) (page 11-25).

See Also:

Oracle Key Vault Administrator's Guide for information about Oracle Key Vault

Changes in Oracle Automatic Storage Management 12c Release 1 (12.1.0.1)

The following are changes in Oracle Automatic Storage Management 12c Release 1 (12.1.0.1).

- [New Features](#) (page xlii)
- [Deprecated Features](#) (page xlv)
- [Desupported Features](#) (page xlv)

New Features

These are new features for Oracle ASM 12c Release 1 (12.1.0.1):

- General Oracle ASM Enhancements

This feature provides general enhancements to Oracle ASM, including:

- Revised version of the physical metadata replication point

Oracle ASM now replicates physically addressed metadata, such as the disk header and allocation tables, within each disk. This enhancement ensures that Oracle ASM is more resilient to bad disk sectors and external corruptions. The disk group attribute `PHYS_META_REPLICATED` is provided to track the replication status of a disk group.

For more information on the `PHYS_META_REPLICATED` attribute, refer to [PHYS_META_REPLICATED](#) (page 4-7).

- Support for increased storage limits

Oracle ASM now supports 511 disk groups. The maximum Oracle ASM disk size is now 32 petabytes (PB).

For information about Oracle ASM size limits, refer to [Oracle ASM Storage Limits](#) (page 4-49).

- The `ALTER DISKGROUP` statement has been updated with a `REPLACE` clause.

For more information about replacing disks in a disk group, refer to [Replacing Disks in Disk Groups](#) (page 4-24).

- Oracle Flex ASM

Oracle Flex ASM enables an Oracle ASM instance to run on a separate physical server from the database servers. Many Oracle ASM instances can be clustered to support a large number of database clients.

This feature enables you to consolidate all the storage requirements into a single set of disk groups. All these disk groups can be managed by a small set of Oracle ASM instances running in a single cluster.

Oracle Flex ASM supports only Oracle Database 12c Release 1 (12.1) and later releases.

For information, refer to [Managing Oracle Flex ASM](#) (page 19-1).

- Oracle ASM Disk Scrubbing

Oracle ASM disk scrubbing checks logical data corruptions and repairs the corruptions automatically in normal and high redundancy disks groups. The feature is designed so that it does not have any impact to the regular input and output (I/O) operations in production systems. The scrubbing process repairs logical corruptions using the Oracle ASM mirror disks. Disk scrubbing uses Oracle ASM rebalancing to minimize I/O overhead.

The scrubbing process is visible in fields of the V\$ASM_OPERATION view.

For more information, refer to [Scrubbing Disk Groups](#) (page 4-32).

- Oracle ASM Disk Resync Enhancements

The disk resync enhancements enable fast recovery from instance failure and faster resync performance overall. Oracle ASM disk resync enables multiple disks to be brought online simultaneously or to control the speed of the resync operation. Oracle ASM disk resync has a resync power limit to control resync parallelism and improve performance. Disk resync checkpoint functionality provides faster recovery from instance failures by enabling the resync to resume from the point at which the process was interrupted or stopped, instead of starting from the beginning.

For more information about disk resync, refer to [Oracle ASM Fast Mirror Resync](#) (page 4-43).

- Oracle ASM Rebalance Enhancements

Oracle ASM rebalance enhancements improve scalability, performance, and reliability of the rebalance operation. This feature extends the rebalance process to operate on multiple disk groups in a single instance. In addition, this feature improves support for thin provisioning, user-data validation, and improved error handling.

For more information about disk group rebalancing, refer to [Manually Rebalancing Disk Groups](#) (page 4-28) and [Tuning Rebalance Operations](#) (page 4-31).

- Shared Oracle ASM Password File in a Disk Group

This feature implements the infrastructure needed to address the bootstrapping issues of Oracle ASM shared password file in an Oracle ASM disk group. This feature solves the bootstrapping problem for storing shared Oracle ASM password files in a disk group.

For information about a shared password file in a disk group, refer to [Managing a Shared Password File in a Disk Group](#) (page 3-33). For information about ASMCMD commands to manage an Oracle ASM or database instance password file in a disk group, refer to [ASMCMD Instance Management Commands](#) (page 10-11).

See Also:

- *Oracle Database Administrator's Guide* for information about ORAPWD commands
 - *Oracle Database Administrator's Guide* for information SRVCTL commands
 - *Oracle Real Application Clusters Administration and Deployment Guide* for information about SRVCTL commands
-
-

- Oracle ASM File Access Control Enhancements on Windows

This feature provides enhancements to Oracle ASM File Access Control on Windows platforms.

Oracle Database now supports the use of an Oracle home user, which can be specified at installation time. The Oracle home user is associated with an Oracle home and it cannot be changed after installation. Different Oracle homes on a system can share the same Oracle home user or use different Oracle home user names. For more information, refer to [About Oracle Home User on Windows](#) (page 3-24).

You can now use access control to separate roles in Windows environments. With Oracle Database services running as users rather than Local System, the Oracle ASM access control feature is enabled to support role separation on Windows. In previous releases, this feature was disabled on Windows because all Oracle services run as Local System.

You can change the identity of an Oracle ASM user from one operating system (OS) user to another OS user without having to drop and re-create the user, which requires dropping all the files a user owns, which improves the manageability of Oracle ASM users and the files they own.

You can modify Windows file access controls while files are open using ASMCMD file access control commands, such as `chgrp`, `chmod`, and `chown`.

For more information about Oracle ASM File Access Control, refer to [Managing Oracle ASM File Access Control for Disk Groups](#) (page 4-61). For information about Oracle ASM File Access Control ASMCMD commands, refer to [ASMCMD File Access Control Commands](#) (page 10-90).

See Also:

Oracle Database Platform Guide for Microsoft Windows for information about running Oracle services on Windows platforms, different types of Windows user accounts, and access control on Windows platforms

- Rolling Migration Framework for Oracle ASM One-off Patches

This feature enhances the rolling migration framework to apply oneoff patches released for Oracle ASM in a rolling manner.

You can use this feature in a clustered Oracle ASM 12c Release 1 (12.1) and higher environment to update one node at a time to the latest patch level without affecting the overall availability of the Oracle ASM cluster or the database clusters using Oracle ASM for storage. This feature improves database availability by

migrating the database to another Oracle ASM instance before a shut down and upgrade operation.

For information, refer to [About Oracle ASM Rolling Patches](#) (page 3-26).

See Also:

Oracle Grid Infrastructure Installation Guide for information about rolling migration of Oracle ASM and Oracle Grid Infrastructure homes

- Updated Key Management Framework

This feature updates Oracle key management commands to unify the key management application programming interface (API) layer. The updated key management framework makes interacting with keys in the wallet easier and adds new key metadata that describes how the keys are being used.

This feature adds a new capability to store a credential file (wallet) in Oracle ASM disk groups. For information about file types supported by Oracle ASM, refer to [What Types of Files Does Oracle ASM Support?](#) (page 5-1).

See Also:

Oracle Database Advanced Security Guide for information about the Oracle key management framework

- Oracle Enterprise Manager

Oracle Enterprise Manager replaces Enterprise Manager Database Control for administering Oracle Automatic Storage Management. For more information, refer to [Administering Oracle ASM with Oracle Enterprise Manager](#) (page 7-1).

Deprecated Features

These are deprecated features for Oracle ASM 12c Release 1 (12.1.0.1).

- Deprecation of single letter SRVCTL CLI options in future releases

All SRVCTL commands have been enhanced to accept more meaningful and user-friendly keyword-style options instead of the single letter options. All new SRVCTL command line interface options added in this release support keyword-style options only and do not have single-letter equivalents. The previous style of using single-letter keyword options as existed in previous releases remain fully supported. In a future release, the single-letter options may no longer be supported and customers may want to change their scripts to adapt the keyword-style command options.

Desupported Features

These are desupported features for Oracle ASM 12c Release 1 (12.1.0.1).

- Desupport of RAW (Block) Storage Devices for Oracle Database and Related Technologies

Oracle Database 12c and related technologies, such as Oracle Clusterware 12c, no longer support the use of raw (block) storage devices. This means that files must be moved to Oracle ASM before upgrading to Oracle Clusterware 12c.

Changes in Oracle ACFS and Oracle ADVM for Oracle ASM 12c Release 1 (12.1.0.1)

The following are changes for Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) in Oracle ASM 12c Release 1 (12.1.0.1).

New Features

These are new features for Oracle ACFS and Oracle ADVM in Oracle ASM 12c Release 1 (12.1.0.1):

- Oracle ACFS Enhancements

This feature introduces enhancements for Oracle ACFS, including:

- Oracle ACFS in Oracle Grid 12c Release 1 (12.1) supports all database files for Oracle Database 11g Release 2 (11.2.0.4) or later, except for data files and redo logs in an Oracle Restart (standalone server) configuration. For this feature, the `COMPATIBLE .ADVM` attribute must be set to 12.1 or later for the disk group that contains the Oracle ACFS file system.

Using Oracle ACFS replication or encryption with database files on Oracle ACFS is not supported. For information about other replication options for database files on Oracle ACFS, refer to *Oracle Data Guard Concepts and Administration* and Oracle GoldenGate documentation. Oracle GoldenGate is an Oracle product sold independently of the Oracle Database. To encrypt database data files on Oracle ACFS, Oracle recommends Oracle Advanced Security. Oracle Advanced Security provides Transparent Data Encryption (TDE) to encrypt data files for entire tablespaces. For information about Transparent Data Encryption (TDE), refer to *Oracle Database Advanced Security Guide*.

- Oracle ACFS supports clusterwide, file granular `fcntl` advisory file locking while byte range locks continue to operate in node local mode.
- Oracle ACFS directory listing performance is improved for newly created directories.
- Oracle ACFS supports end-to-end storage visibility for files with the `-d` option of `acfsutil info file`. The `-d` option displays details about a file extent location, from the Oracle ACFS extent to the Oracle ASM devices in the disk group. For information, refer to [acfsutil info file](#) (page 16-147).
- Oracle ACFS ASMCMD commands have enhanced syntax. For information about ASMCMD commands for Oracle ACFS, refer to [ASMCMD Volume Management Commands](#) (page 15-1).
- Oracle ACFS supports unlimited expansions when resizing a file system in a disk group with ADVM compatibility set to 11.2.0.4 or higher.
- Oracle ACFS performance and scalability improvements for `ls` and `find` in a disk group with ADVM compatibility set to 11.2.0.4 or higher.

For more information about Oracle ACFS, refer to [Overview of Oracle ACFS](#) (page 11-1)

- Oracle ACFS Enhancements Supporting Network File Storage (NFS)

Note:

This feature is not supported in Oracle Restart configurations.

High Availability NFS for Oracle Grid Infrastructure provides uninterrupted service of NFS V2/V3/V4 exported paths by exposing NFS exports on Highly Available Virtual IPs (HAVIP) and using Oracle Clusterware agents to ensure that the HAVIPs and NFS exports are always online. While base NFS supports file locking, HANFS does not support NFS file locking.

Oracle ACFS file system exports (using NFS) include Golden Images and patch updates applied to Oracle ACFS snapshots. Note that NFS is deployed with Oracle Grid Infrastructure servers in support of Oracle Grid Infrastructure client systems. Application VIP and NFS export resources are employed for Oracle ACFS and highly available NFS.

High Availability NFS can also be used as a general high availability solution for Oracle ACFS exported file systems.

For information about High Availability NFS for Oracle Grid Infrastructure homes, refer to [Oracle ACFS NAS Maximum Availability eXtensions](#) (page 11-44).

See Also:

- *Oracle Real Application Clusters Administration and Deployment Guide* for information about SRVCTL commands
 - *Oracle Clusterware Administration and Deployment Guide* for information about Oracle Clusterware and clusters
-
-

- Enhancements for Oracle ACFS Snapshots

Oracle ACFS supports the creation of a snapshot from an existing snapshot in the same Oracle ACFS file system. In addition, snapshot conversions are enabled between read-only and read-write snapshots.

For information about Oracle ACFS snapshots, refer to [About Oracle ACFS Snapshots](#) (page 11-11). For information about Oracle ACFS snapshot commands, refer to [Oracle ACFS Command-Line Utilities](#) (page 16-143).

- Oracle ACFS Tagging Generic Application Programming Interface

Tagging operations are available as C application programming interfaces (APIs), described in [Oracle ACFS Tagging Generic Application Programming Interface](#) (page 17-12).

For information about Oracle ACFS tagging, refer to [Oracle ACFS Tagging](#) (page 11-39). For information about Oracle ACFS tagging operations available through the `acfsutil tag` commands, refer to [Oracle ACFS Command-Line Tools for Tagging](#) (page 16-60).

- Oracle Clusterware Resources Supporting Oracle ACFS

Note:

This feature is not supported in Oracle Restart configurations.

Oracle ACFS and Oracle ADVM resource support is enhanced for both performance and coverage. Oracle Clusterware (CRS) managed resource support is extended to include enhancements for Oracle ACFS Oracle homes, Oracle ACFS General Purpose file systems, and Oracle ADVM volumes in supporting Oracle Grid Infrastructure Server Oracle ACFS export (using NFS) configurations. SRVCTL has updates to commands for managing Oracle ACFS and Oracle ADVM resources.

For information about Oracle Clusterware resources and Oracle ACFS, refer to [Oracle Clusterware Resources and Oracle ACFS Administration](#) (page 11-18).

See Also:

- *Oracle Real Application Clusters Administration and Deployment Guide* for information about SRVCTL commands
 - *Oracle Clusterware Administration and Deployment Guide* for information about Oracle Clusterware resources
-
-

- Oracle ACFS Plugins

The Oracle ACFS plugin functionality enables a user space application to collect *just-in-time* Oracle ACFS file and Oracle ADVM volume metrics from the operating system environment. Applications can use the Oracle ACFS plug-in infrastructure to create customized solutions that extend the general application file metric interfaces to include detailed Oracle ACFS file system and volume data.

For information about Oracle ACFS plugins, refer to [Oracle ACFS Plugins](#) (page 11-43). For information about Oracle ACFS plugin commands, refer to [Oracle ACFS Command-Line Utilities](#) (page 16-143). For information about the Oracle ACFS plug-in application programming interface, refer to [Oracle ACFS Plug-in Generic Application Programming Interface](#) (page 17-8).

- Integration of Oracle ACFS Replication With Auditing, Encryption, and Security

This feature enables the integration of Oracle Automatic Storage Management Cluster File System (Oracle ACFS) replication with Oracle ACFS auditing, encryption, and security.

For more information, refer to [Using Replication with Auditing, Encryption, and Security](#) (page 11-41).

- Oracle ACFS Auditing and Support for Importing Auditing Data into Audit Vault Server

This feature provides auditing for Oracle ACFS security and encryption. In addition, this feature also generates an XML file containing Oracle ACFS audit trail data which can be imported by Audit Vault Server.

For more information, refer to [Oracle ACFS Auditing](#) (page 11-29). For information about Oracle ACFS commands for auditing, refer to [Oracle ACFS Command-Line Tools for Auditing](#) (page 16-81) and the `acfsutil sec realm enable` and

`acfsutil sec realm disable` commands described in [Oracle ACFS Command-Line Tools for Security](#) (page 16-86).

- Support for Oracle ACFS Replication and Tagging on Solaris
This feature supports Oracle ACFS replication and tagging on Solaris platforms. For information about Oracle ACFS, refer to [Introducing Oracle ACFS and Oracle ADVM](#) (page 11-1).
- Support for Oracle ACFS Replication and Tagging on AIX
This feature supports Oracle ACFS replication and tagging on AIX platforms. For information about Oracle ACFS, refer to [Introducing Oracle ACFS and Oracle ADVM](#) (page 11-1).
- Support for Oracle ACFS Security and Encryption on Solaris
This feature extends support for Oracle ACFS Security and Encryption Features on Solaris systems. Changes in functionality for Solaris environments are noted with the individual security and encryption commands. For more information, refer to [Oracle ACFS Security](#) (page 11-22) and [Oracle ACFS Encryption](#) (page 11-25).
- Oracle Enterprise Manager
Oracle Enterprise Manager replaces Enterprise Manager Database Control for administering Oracle Automatic Storage Management Cluster File System. For more information, refer to [Administering Oracle ACFS with Oracle Enterprise Manager](#) (page 13-1).

Part I

Introduction

Part I provides an overview of Oracle Automatic Storage Management (Oracle ASM), Oracle ASM Cluster File System (Oracle ACFS), and Oracle ASM Dynamic Volume Manager (Oracle ADVM) concepts and features in the following chapters:

- [Introducing Oracle Automatic Storage Management](#) (page 1-1)
- [Exploring Considerations for Oracle ASM Storage](#) (page 2-1)

Introducing Oracle Automatic Storage Management

Oracle Automatic Storage Management (Oracle ASM), Oracle ASM Cluster File System (Oracle ACFS), and Oracle ASM Dynamic Volume Manager (Oracle ADVM) are key components of storage management.

This chapter provides an overview of Oracle Automatic Storage Management (Oracle ASM), Oracle ASM Cluster File System (Oracle ACFS), and Oracle ASM Dynamic Volume Manager (Oracle ADVM) concepts and features. This chapter contains the following topics:

- [Overview of Oracle Automatic Storage Management](#) (page 1-1)
- [About Oracle Automatic Storage Management Cluster File System](#) (page 1-2)
- [Understanding Oracle ASM Concepts](#) (page 1-3)
- [Understanding Oracle ASM Disk Group Administration](#) (page 1-11)

For a list of the terms that are used in the *Oracle Automatic Storage Management Administrator's Guide* and their definitions, refer to the [Glossary](#) (page 1) in this guide.

See Also:

- The Oracle Cloud Storage page on the Oracle Technology Network website at <http://www.oracle.com/technetwork/database/cloud-storage/index.html> for more information about Oracle ASM
 - "Master Note for Automatic Storage Management (ASM) (Doc ID 1187723.1)" at [My Oracle Support](https://support.oracle.com) (<https://support.oracle.com>)
-

1.1 Overview of Oracle Automatic Storage Management

Oracle ASM is Oracle's recommended storage management solution that provides an alternative to conventional volume managers, file systems, and raw devices.

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 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.

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 ASM has easy to use management interfaces such as SQL*Plus, the Oracle ASM Command Line Utility (ASMCMD) command-line interface, and Oracle ASM Configuration Assistant (ASMCA).

See Also:

- [Administering Oracle ASM Disk Groups](#) (page 4-1) for information about administering disk groups
 - [Managing Oracle ASM With ASMCA](#) (page 9-1) for information about Oracle ASM Configuration Assistant
 - [Managing Oracle ASM with ASMCMD](#) (page 10-1) for information about the ASMCMD command-line interface
 - *Oracle Database Administrator's Guide* for information about Oracle Database structure and storage
-
-

1.2 About Oracle Automatic Storage Management Cluster File System

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) extend Oracle ASM functionality.

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 all customer files. The Oracle ASM Dynamic Volume Manager (Oracle ADVM) provides volume management services and a standard disk device driver interface to clients.

See Also:

[Oracle Automatic Storage Management Cluster File System](#) (page 1) for more information about Oracle ACFS and Oracle ADVM.

1.3 Understanding Oracle ASM Concepts

The concepts for the key Oracle ASM components are introduced in this topic.

The following topics are discussed:

- [About Oracle ASM Instances](#) (page 1-3)
- [About Oracle ASM Disk Groups](#) (page 1-6)
- [About Mirroring and Failure Groups](#) (page 1-6)
- [About Oracle ASM Disks](#) (page 1-7)
- [About Oracle ASM Allocation Units](#) (page 1-7)
- [About Oracle ASM Files](#) (page 1-8)

See Also:

[Exploring Considerations for Oracle ASM Storage](#) (page 2-1) for information about preparing your storage environment.

1.3.1 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 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 file names 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 in a Standard Oracle ASM cluster fails, then all of the database instances on that node also fail. However, in an Oracle Flex ASM configuration, Oracle 12c database instances would not fail as they would be able to access another Oracle ASM instance remotely on another node.

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 (page 1-4) 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.

In Figure 1-1 (page 1-4), 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 (page 1-4) shows multiple database instances, but only one Oracle ASM instance is needed to serve the multiple database instances.

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

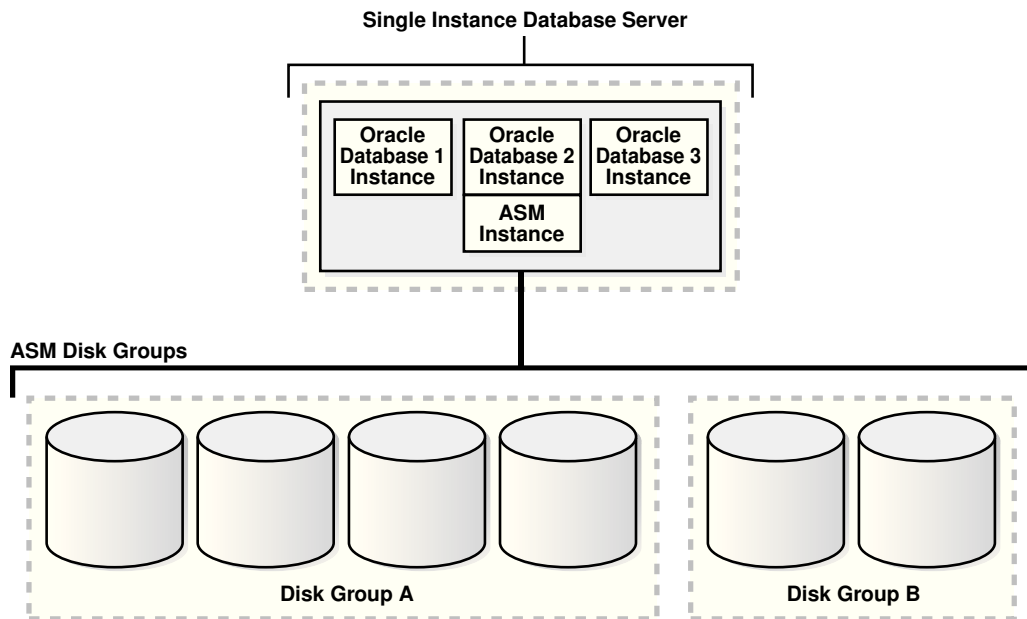



Figure 1-2 (page 1-5) shows an Oracle ASM cluster in an Oracle RAC environment where Oracle ASM provides a clustered pool of storage. There is one Oracle ASM instance for each node serving multiple Oracle RAC or single-instance databases in the cluster. All of the databases are consolidated and share the same two Oracle ASM disk groups.

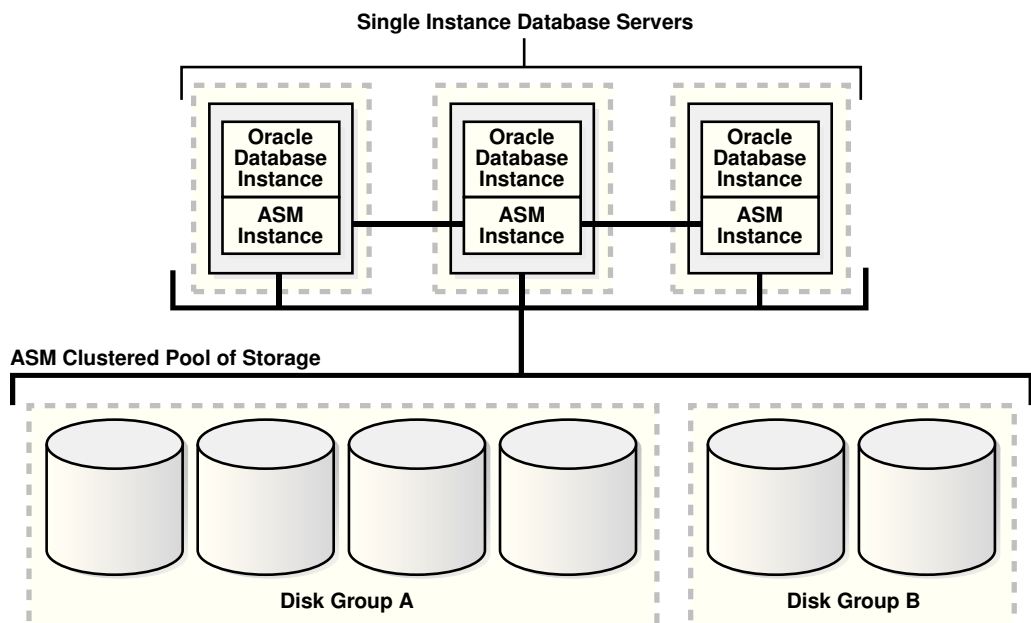
Figure 1-2 Oracle ASM Cluster Configuration with Oracle RAC



A clustered storage pool can be shared by multiple single-instance Oracle Databases as shown in [Figure 1-3](#) (page 1-5). 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



See Also:

- [Administering Oracle ASM Instances](#) (page 3-1) for information about managing an Oracle ASM instance.
 - [Managing Oracle Flex ASM](#) (page 19-1) for information about Oracle Flex ASM configurations.
-

1.3.2 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.

See Also: [Administering Oracle ASM Disk Groups](#) (page 4-1) for more information about managing disk groups

1.3.3 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.

Oracle ASM mirroring is more flexible than traditional RAID mirroring. For a disk group specified as `NORMAL` redundancy, 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.

See Also:

- [Managing Disk Group Templates](#) (page 5-19) for information about disk group types and templates
 - [Mirroring, Redundancy, and Failure Group Options](#) (page 4-38) for more information about mirroring and failure groups
-

1.3.4 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 path name 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.

1.3.5 About Oracle ASM 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.

See Also:

- [Example 4-1](#) (page 4-13) for an example that shows how the `AU_SIZE` is specified with the `CREATE DISKGROUP SQL` statement
 - [Features Enabled By Disk Group Compatibility Attribute Settings](#) (page 4-57) for information about allocation unit sizes and disk group compatibility attributes
 - Articles at [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com) for additional information about allocation units and Oracle ASM
-

1.3.6 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 topics describe Oracle ASM file components:

- [Understanding Oracle ASM Extents](#) (page 1-8)
- [Understanding Oracle ASM Striping](#) (page 1-9)
- [About Oracle ASM File Templates](#) (page 1-11)

1.3.6.1 Understanding Oracle ASM 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. The various extent sizes are described in this topic.

For disk groups with AU size less than 4 MB:


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

For disk groups with AU size greater than or equal to 4 MB and the disk group RDBMS compatibility greater than or equal to 11.2.0.4, the counts for extents of sizes (the disk group AU size, 4*AU size, or 16*AU size) are calculated using the application block size to support maximum file size.

The extent sizing 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](#) (page 4-50)".

[Figure 1-4](#) (page 1-9) 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 the 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](#) (page 1-9)).

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



1.3.6.2 Understanding Oracle ASM Striping

Oracle ASM striping has two primary purposes: balance loads across all of the disks in a disk group and 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 fine-grained stripe size always equals 128 KB in any configuration; this provides lower I/O latency for small I/O operations. The coarse-grained stripe size is always equal to the AU size (not the data extent size).

Figure 1-5 (page 1-10) and Figure 1-6 (page 1-11) are illustrations of Oracle ASM file striping. In both illustrations, the allocation unit size has been set to 1 M (AU_SIZE = 1M) for the disk group which consists of 8 disks. The instance is Oracle ASM 11g Release 2 (11.2) and the disk group compatibility attributes for ASM and RDBMS have been set to 11.2, so variable extents are shown in the graphic after the first 20,000 extents. For the first 20,000 extents, the extent size is 1 M and equals one allocation unit (AU). For the next 20,000 extents, the extent size is 4 M and equals 4 AUs.


To identify the stripe chunks of the file, they have been labeled A..X (24 letters) using different fonts for successive series of A..X until all the chunks have been identified.

In Figure 1-5 (page 1-10), the file is striped in 128 K chunks (labeled A..X) with each 128 K chunk stored in an extent, starting at the first extent in disk 1, then the first extent in disk 2, and then continuing in a round-robin pattern through all the disks until the entire file has been striped. As shown in this example, the striping chunks first fill up the first extent of each disk, then the second extent of each disk, and so on until the entire file has been striped.

Figure 1-5 Oracle ASM Fine-Grained Striping



In Figure 1-6 (page 1-11), the file is striped in 1 M chunks (labeled A..X) with each 1 M chunk stored uniquely in an extent, starting at the first extent in disk 1, then the first extent in disk 2, and then continuing in a round-robin pattern through all the disks until the entire file has been striped. For the first 20,000 extents where the AU equals the extent size (1 M), the stripe equals the extent size and allocation unit size. For the variable extents, where an extent is composed of multiple allocation units, the file stripe is located in an AU of the extent. The striping chunks are placed in the allocation units of the first extents of all the disks before the striping continues to the next extent.

Figure 1-6 Oracle ASM Coarse-Grained Striping

1.3.6.3 About Oracle ASM 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.

See Also:

[Managing Disk Group Templates](#) (page 5-19) for more information about Oracle ASM templates

1.4 Understanding Oracle ASM Disk Group Administration

Oracle ASM disk group administration is introduced in this topic.

The following topics are discussed:

- [About Discovering Disks](#) (page 1-11)
- [About Mounting and Dismounting Disk Groups](#) (page 1-12)
- [About Adding and Dropping Disks](#) (page 1-12)
- [About Online Storage Reconfigurations and Dynamic Rebalancing](#) (page 1-13)

1.4.1 About Discovering Disks

The disk discovery process locates the operating system names for disks that Oracle ASM can access.

Disk discovery finds all of the disks that comprise a disk group to be mounted. The set of discovered disks also includes disks that could be added to a disk group.

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

See Also:

- [ASM_DISKSTRING](#) (page 3-10) for information about the `ASM_DISKSTRING` initialization parameter
 - [Oracle ASM Disk Discovery](#) (page 4-33) for information about disk discovery
 - [Administering Oracle ASM Filter Driver](#) (page 20-1) for information about Oracle ASM Filter Driver
-
-

1.4.2 About Mounting and Dismounting 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.

See Also:

[Mounting and Dismounting Disk Groups](#) (page 4-67) for more information about disk groups

1.4.3 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 could 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.

See Also:

[Altering Disk Groups](#) (page 4-19) for more information about altering disk group membership

1.4.4 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 based on I/O 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 `ASM_POWER_LIMIT` initialization parameter.

See Also:

- [ASM_POWER_LIMIT](#) (page 3-11) for more information about the power limit setting
 - [Manually Rebalancing Disk Groups](#) (page 4-28) for more information about disk rebalancing
-

Exploring Considerations for Oracle ASM Storage

Several issues should be considered 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 must also scale up the storage subsystem.

This chapter contains the following topics:

- [Storage Resources for Disk Groups](#) (page 2-1)
- [Oracle ASM and Multipathing](#) (page 2-3)
- [Recommendations for Storage Preparation](#) (page 2-4)
- [Storage Device Path and Permission Persistence](#) (page 2-6)

2.1 Storage Resources for Disk Groups

There are various storage resources that can be used to create Oracle ASM 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 provides mirroring.

- **Network File System (NFS)**

An Oracle ASM disk group can be created from NFS files, including Oracle Direct NFS (dNFS). The NFS files that are provisioned to a disk group may be from multiple NFS servers to provide better load balancing and flexible capacity planning.

You can use NFS, with or without Direct NFS, to store data files. However, NFS is not supported for Oracle Clusterware files. To install Oracle Real Application Clusters (Oracle RAC) on Windows using NFS, you must also have access to a shared storage method other than NFS for Oracle Clusterware files.

NFS-based quorum disks (quorum failure groups) should not use Direct NFS (dNFS) because dNFS does not support soft mounts. Instead, use a soft mount of a NFS mount point for quorum disks. When using a soft mount, Oracle ASM handles an I/O failure gracefully, and sets only the associated quorum disk offline.

With hard mounts, the Oracle ASM or the database instance may hang if the NFS server becomes unavailable. Note that these hang situations can occur whether or not Direct NFS is used and whether or not Oracle ASM is used for mirroring.

See Also:

Oracle Grid Infrastructure Installation and Upgrade Guide for your operating system for information about Oracle Direct NFS and storage requirements for Oracle ASM

Note:

- Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes and Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems are currently not supported on disk groups that have been created from NFS or Common Internet File System (CIFS) files. However, Oracle ACFS file systems may be exported as NFS or CIFS file systems to network clients in some cases. Samba/CIFS clients on Windows cannot use ACLs when interfacing with Oracle ACFS Linux, Solaris, or AIX servers.
 - Mounting loopback file systems over Oracle ACFS files is not supported.
 - Block or raw devices are not supported by Oracle Universal Installer (OUI) or Database Configuration Assistant (DBCA).
-
-

The procedures for preparing storage resources for Oracle ASM are:

- 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 device names are typically presented from the `/dev` directory with the `/dev/device_name_identifier` name syntax.

- 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](#) (page 3-28).

- Change the device permissions to read/write

Note:

To ensure that ownership and permission settings are persistent, you can use `udev` to ensure that the disks do not revert to root ownership when the systems restart.

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](#) (page 3-10).

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 software. 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](#) (page 3-28).

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

2.2 Oracle ASM and Multipathing

Multipathing solutions provide failover by using redundant physical path components.

These redundant physical path 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 then 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.

See Also:

- [Oracle ASM Disk Discovery](#) (page 4-33) for information about disk discovery
 - [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com) for articles with information about Oracle ASM and Multipathing
 - Your platform-specific installation guide for information about configuring multipathing for your system
-
-

2.3 Recommendations for Storage Preparation

Recommendations for storage preparation with Oracle ASM are discussed in this topic.

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

- Configure a separate disk group for the following:
 - Oracle Cluster Registry (OCR) and voting files
 - Grid Infrastructure Management Repository (GIMR) files
 - Database data files
 - Fast recovery area

- The number of LUNs (Oracle ASM disks) for each disk group should be at least equal to four times the number of active I/O paths. For example, if a disk group has two active I/O paths, then minimum of eight LUNs should be used. The LUNs should be of equal size and performance for each disk group.

An I/O path is a distinct channel or connection between storage presenting LUNs and the server. An *active* I/O path is an I/O path in which the I/O load on a LUN is multiplexed through multipathing software.

- Ensure that all Oracle ASM disks in a disk group have similar storage performance and availability characteristics. In storage configurations with mixed speed drives, such as flash memory and hard disk drives (HDD), 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.
- Configure a minimum of three failure groups for normal redundancy disk groups and five failure groups for high redundancy disk groups to maintain the necessary number of copies of the Partner Status Table (PST) to ensure robustness with respect to storage hardware failures.
- 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 when not using 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.
- Use the Oracle ASM Filter Driver feature to provide consistent device naming and permission persistency.

See Also:

- [Creating Disk Groups for a New Oracle Installation](#) (page 4-14) for information about configuring disk groups for a new installation
 - The Oracle Cloud Storage page on the Oracle Technology Network website at <http://www.oracle.com/technetwork/database/cloud-storage/index.html> for more information about Oracle ASM
 - *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
 - [Oracle ASM Failure Groups](#) (page 4-40) for information about failure groups
 - [Administering Oracle ASM Filter Driver](#) (page 20-1) for Information about Oracle ASM Filter Driver
-
-

2.4 Storage Device Path and Permission Persistence

Before installation, or before configuring new storage devices to use with Oracle ASM, administrators must configure storage device names and ownership to ensure that storage paths and ownership persist after system restarts.

Use Oracle ASM Filter Driver to maintain permissions and manage device paths. On Oracle Solaris, you can also use the Solaris I/O multipathing features to maintain permissions and device paths.

See Also:

Oracle Grid Infrastructure Installation and Upgrade Guide for your operating system for more information about configuring storage devices for path and permission persistence.

Part II

Oracle ASM Instances and Disk Groups

Part II discusses the administration of Oracle ASM instances and disk groups in the following chapters:

- [Administering Oracle ASM Instances](#) (page 3-1)
- [Administering Oracle ASM Disk Groups](#) (page 4-1)
- [Administering Oracle ASM Files, Directories, and Templates](#) (page 5-1)
- [Using Views to Display Oracle ASM Information](#) (page 6-1)
- [Administering Oracle ASM with Oracle Enterprise Manager](#) (page 7-1)
- [Performing Oracle ASM Data Migration with RMAN](#) (page 8-1)
- [Managing Oracle ASM With ASMCA](#) (page 9-1)
- [Managing Oracle ASM with ASMCMD](#) (page 10-1)

Administering Oracle ASM Instances

Administering Automatic Storage Management (Oracle ASM) instances is similar to administering Oracle Databases.

This chapter describes how to administer Automatic Storage Management (Oracle ASM) instances, configure Oracle ASM instance parameters, and 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 ASM Command Line Utility (ASMCMD) command-line interface, Oracle ASM Configuration Assistant (ASMCA), and SQL*Plus to perform Oracle ASM instance administration tasks.

This chapter contains the following topics:

- [Operating with Different Releases of Oracle ASM and Database Instances Simultaneously](#) (page 3-1)
- [Initialization Parameters for Oracle ASM Instances](#) (page 3-2)
- [Administration of Oracle ASM Instances](#) (page 3-17)
- [Upgrading and Patching Oracle ASM](#) (page 3-25)
- [Authentication for Accessing Oracle ASM Instances](#) (page 3-27)
- [Administering the Oracle ASM Audit Trail](#) (page 3-36)
- [Migration of a Database to Use Oracle ASM](#) (page 3-37)

See Also:

- [About Oracle ASM Instances](#) (page 1-3) for a description of an Oracle ASM instance
 - [Administering Oracle ASM Filter Driver](#) (page 20-1) for information about Oracle ASM Filter Driver
-

3.1 Operating with Different Releases of Oracle ASM and Database Instances Simultaneously

Oracle Automatic Storage Management (Oracle ASM) in Oracle Database 12c Release 2 (12.2) supports Oracle Database 12c Release 2 (12.2) or older software versions, including Oracle Database 10g Release 1 (10.1).

Note:

- An Oracle ASM instance must be at Oracle ASM 12c Release 2 (12.2) to support Oracle Database 12c Release 2 (12.2).
 - Refer to *Oracle Exadata documentation* for information about the Oracle Database versions that Oracle ASM 12c Release 2 (12.2) 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.
-

When using different software versions, the database instance supports Oracle ASM functionality of the earliest release in use. For example, an Oracle Database 10g Release 1 (10.1) database instance operating with an Oracle ASM 12c Release 2 (12.2) instance only supports Oracle ASM 10g Release 1 (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 the `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.

See Also:

- [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility attributes settings
 - [Example 6-4](#) (page 6-4) for an example showing a query on the `V$ASM_CLIENT` view
 - [Views Containing Oracle ASM Disk Group Information](#) (page 6-1) for more information about the `V$ASM_CLIENT` and `V$ASM_*` views
-

3.2 Initialization Parameters for Oracle ASM Instances

Settings for Oracle ASM instance initialization parameters are discussed in this topic.

To install and initially configure an Oracle ASM instance, use Oracle Universal Installer (OUI) and Oracle ASM Configuration Assistant (ASMCA).

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. Only a few Oracle ASM-specific instance initialization parameters must be configured. The default values are usually sufficient.

The following topics are discussed:

- [About Initialization Parameter Files for an Oracle ASM Instance](#) (page 3-3)
- [About Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File](#) (page 3-4)
- [About Oracle ASM Initialization Parameter Settings](#) (page 3-7)
- [About Automatic Memory Management for Oracle ASM](#) (page 3-7)
- [Recommended Settings for Oracle ASM Initialization Parameters](#) (page 3-8)
- [About Setting Database Initialization Parameters for Use with Oracle ASM](#) (page 3-16)

See Also:

- *Oracle Database Reference* for information about initialization parameters
- *Oracle Database Administrator's Guide* for information about initialization parameter files
- *Oracle Grid Infrastructure Installation and Upgrade Guide* for information about installing and configuring Oracle ASM
- The Oracle Cloud Storage page on the Oracle Technology Network website at <http://www.oracle.com/technetwork/database/cloud-storage/index.html> for more information about Oracle ASM best practices

3.2.1 About Initialization Parameter Files for an Oracle ASM Instance

When installing Oracle ASM in an Oracle Restart (standalone) configuration, Oracle Universal Installer (OUI) creates a separate server parameter file (SPFILE) and password file for the Oracle ASM instance. The ASM SPFILE is stored in a disk group during installation.

When installing Oracle ASM in a clustered Oracle ASM environment, OUI creates a single, shared SPFILE for Oracle ASM in a disk group.

When upgrading an Oracle ASM instance, if the ASM SPFILE was originally in a shared file system, then the upgraded Oracle ASM instance retains the SPFILE in the same location. If the original Oracle ASM instance used a PFILE, then after an upgrade the instance continues to use a PFILE.

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 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 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, then 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

Note:

A PFILE or SPFILE is required if your configuration uses nondefault initialization parameters for the Oracle ASM instance.

You can administer Oracle ASM initialization parameter files with SQL*Plus, ASMCA, and ASMCMD commands.

See Also:

- [Managing Oracle ASM With ASMCA](#) (page 9-1) for information about the ASMCA GUI and command-line interfaces
 - [ASMCMD Instance Management Commands](#) (page 10-11) for information about ASMCMD commands for managing an Oracle ASM SPFILE; such as `spbackup`, `spcopy`, and `spmmove`
 - *Oracle Database Administrator's Guide* for more information about creating and maintaining initialization parameter files
 - *Oracle Database SQL Language Reference* for information about creating an SPFILE with the `CREATE SPFILE` SQL statement
-
-

3.2.2 About 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.

In addition, you can use the SQL `CREATE SPFILE` to create an Oracle ASM SPFILE when connected to the Oracle ASM instance.

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

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

The following topics are discussed:

- [Creating, Copying, and Moving an SPFILE Into a Disk Group](#) (page 3-5)

- [Making a Back Up Copy of an Oracle ASM SPFILE in a Disk Group](#) (page 3-6)

See Also:

- [spbackup](#) (page 10-24), [spcopy](#) (page 10-25), and [spmmove](#) (page 10-27) for information about ASMCMD commands for managing an SPFILE
 - *Oracle Database Administrator's Guide* for more information about creating and maintaining 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)
-
-

3.2.2.1 Creating, Copying, and Moving an SPFILE Into a Disk Group

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 instance from Oracle ASM 11g Release 1 (11.1) to Oracle ASM 11g Release 2 (11.2), you could place the Oracle ASM SPFILE in a disk group that has COMPATIBLE .ASM set to 11 . 2. For information about disk group compatibility attributes, see "[Disk Group Compatibility](#) (page 4-50)".

In the following steps, assume an Oracle ASM 11g Release 2 (11.2) instance is using a PFILE stored in \$ORACLE_HOME/dbs/asmpfile.ora. You can use the SQL CREATE SPFILE statement to create an 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 an SPFILE in a disk group, perform the following steps:

1. Connect to the Oracle ASM instance.

For example:

```
$ sqlplus / as sysasm
```

2. Create an 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/asmpfile.ora'
      FROM PFILE = '$ORACLE_HOME/dbs/asmpfile.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](#) (page 10-27)".

Note:

If you include the AS COPY option, then the CREATE SPFILE FROM PFILE statement does not update the clusterwide SPFILE location.

- Restart the Oracle ASM instance so that the instance reads the SPFILE in the new location.

For information on shutting down and starting up an Oracle ASM instance, see "[About Starting Up an Oracle ASM Instance](#) (page 3-18)" and "[About Shutting Down an Oracle ASM Instance](#) (page 3-22)".

3.2.2.2 Making a Back Up Copy of an Oracle ASM SPFILE in a Disk Group

This section describes the steps to make a back up copy of an Oracle ASM SPFILE in another disk group using the ASMCMD commands. If necessary, then the backup copy can be used to restore the Oracle ASM SPFILE.

The source and target disk groups must have the disk group attribute `COMPATIBLE.ASM` set to 11.2 or higher.

To make a copy of the Oracle ASM SPFILE in another disk group with the `spcopy` command perform the following steps:

- Locate the Oracle ASM SPFILE using the ASMCMD `spget` command.

For example:

```
ASMCMD [+] > spget
+DATA/ASM/ASMPARAMETERFILE/registry.253.849343867
```

- Copy the Oracle ASM SPFILE to another disk group with `spcopy` command.

For example:

```
ASMCMD [+] > spcopy +DATA/ASM/ASMPARAMETERFILE/registry.253.849343867 +FRA/spfileCopyASM.ora
```

Running `spcopy` without the `-u` option does not update the location of the Oracle ASM SPFILE. You can use `spset` to set the location of the Oracle ASM SPFILE in the Grid Plug and Play (GPnP) profile.

- List all the copies of the Oracle ASM SPFILE file contained in the FRA disk group using the ASMCMD `ls` command.

For example:

```
ASMCMD [+] > ls -l --absolutepath FRA/ASM/ASMPARAMETERFILE
Type          Redund Striped Time          Sys Name
ASMPARAMETERFILE MIRROR COARSE JUN 06 13:00:00 Y +FRA/spfileCopyASM.ora => REGISTRY.
253.849533009
```

- Verify the current location of the Oracle ASM SPFILE file with the `spget` command.

For example:

```
ASMCMD [+] > spget
+DATA/ASM/ASMPARAMETERFILE/registry.253.849343867
```

In the event that the current Oracle ASM SPFILE file in a disk group has been corrupted or that disk group is not accessible, you can use `spset` or `spcopy` with the `-u` option to restore the Oracle ASM SPFILE file using the backup copy that you have previously created.

For example:

```
ASMCMD [+] > spcopy -u +FRA/spfileCopyASM.ora +DATA2/ASM/spfileASM.ora
```


3.2.3 About Oracle ASM Initialization Parameter Settings

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 parameters after database creation using SQL `ALTER SYSTEM` or `ALTER SESSION` statements.

The Oracle ASM parameters use suitable defaults for most environments. You cannot use parameters with names that are prefixed with `ASM_*` in database instance parameter files.

3.2.4 About 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.

An Oracle ASM instance can automatically increase the values set for `MEMORY_TARGET` and `MEMORY_MAX_TARGET` if an ORA-04031 error is raised and automatic memory management is enabled. If `MEMORY_MAX_TARGET` has been explicitly set to a value, then every time ORA-04031 is raised, the `MEMORY_TARGET` value is increased by 10% of the existing `MEMORY_TARGET` value or 128 MB, whichever is greater, but not greater than the customer specified `MEMORY_MAX_TARGET` value. If `MEMORY_MAX_TARGET` is not explicitly set, then both `MEMORY_TARGET` and `MEMORY_MAX_TARGET` are increased by 10% of the existing `MEMORY_TARGET` value or 128 MB, whichever is greater, for a maximum of five increases. The Oracle ASM instance must be rebooted to use the new `MEMORY_TARGET` and `MEMORY_MAX_TARGET` settings.

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 automatic shared memory management and automatic PGA memory management. 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. Unless specified, the behaviors of the automatic memory management parameters in Oracle ASM instances behave the same as in Oracle Database instances.

Note:

- 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 1 GB. If you set `MEMORY_TARGET` lower, then Oracle increases the value for `MEMORY_TARGET` to 1 GB automatically.
 - For the recommended settings of memory initialization parameters in an Oracle Exadata environment, refer to the Oracle Exadata documentation.
-
-

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
-
-

3.2.5 Recommended Settings for Oracle ASM Initialization Parameters

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

- [ASM_DISKGROUPS](#) (page 3-9)
- [ASM_DISKSTRING](#) (page 3-10)
- [ASM_IO_PROCESSES](#) (page 3-11)
- [ASM_POWER_LIMIT](#) (page 3-11)
- [ASM_PREFERRED_READ_FAILURE_GROUPS](#) (page 3-12)
- [DB_CACHE_SIZE](#) (page 3-13)
- [DIAGNOSTIC_DEST](#) (page 3-14)
- [INSTANCE_TYPE](#) (page 3-14)
- [LARGE_POOL_SIZE](#) (page 3-14)
- [PROCESSES](#) (page 3-15)
- [REMOTE_LOGIN_PASSWORDFILE](#) (page 3-15)
- [SHARED_POOL_SIZE](#) (page 3-16)

See Also:

- *Oracle Database Reference* for more information about initialization parameters
 - *Oracle Database Administrator's Guide* for more information about creating and maintaining an initialization parameter file
-
-

3.2.5.1 ASM_DISKGROUPS

The `ASM_DISKGROUPS` initialization parameter specifies a list of disk group names that an Oracle ASM instance mounts at startup when the `SQL ALTER DISKGROUP ALL MOUNT` statement is issued.

The Oracle ASM instance startup process executes `ALTER DISKGROUP ALL MOUNT` unless the `NOMOUNT` startup option is specified.

The default value of the `ASM_DISKGROUPS` parameter is a `NULL` string.

The `ASM_DISKGROUPS` parameter is dynamic. If you are using a server parameter file (SPFILE), then you do not have to manually alter the value of `ASM_DISKGROUPS`. Oracle ASM automatically adds a disk group to the parameter when the disk group is successfully created or mounted. Oracle ASM also automatically removes a disk group from the 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 Oracle Database 12c Release 1 or later, Oracle ASM configurations support up to 511 disk groups. Oracle ASM configurations with Oracle Database releases before 12c Release 1 can only support up to 63 disk groups.

See Also:

- [About Mounting Disk Groups at Startup](#) (page 3-21) for information about disk groups that are mounted at startup time
 - [Mounting and Dismounting Disk Groups](#) (page 4-67) for additional information about mounting Oracle ASM disk groups
 - *Oracle Database Reference* for more information about the `ASM_DISKGROUPS` initialization parameter
-

3.2.5.2 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, to limit the discovery process to only include disks that are in the `/dev/rdisk/mydisks` directory for an Oracle ASM instance that does not use Oracle ASM Filter Driver (Oracle ASMFDF) or ASMLIB, 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`, you could set `ASM_DISKSTRING` as follows on a Linux system:

```
ASM_DISKSTRING = '/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 have 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.

See Also:

- [Oracle ASM Disk Discovery](#) (page 4-33) for more information about discovering disks
 - *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
-
-

3.2.5.3 ASM_IO_PROCESSES

The `ASM_IO_PROCESSES` initialization parameter specifies the number of IO worker processes to be started in an Oracle IO Server (IOS) server instance.

This parameter is only applicable for an Oracle IO Server instance, which runs in an Oracle Grid Infrastructure home.

IO Server instances are started automatically in the Oracle Grid Infrastructure configuration when generating the credential file for a client cluster. You can configure the number of IO Server instances running in the client cluster.

The default value should work in most cases. However, under heavy IO loads a higher value than the default may be appropriate.

See Also:

- [Overview of Oracle Flex ASM](#) (page 19-1) for information about Oracle IO Server and Oracle Flex ASM
 - *Oracle Database Reference* for more information about the `ASM_IO_PROCESSES` initialization parameter
-
-

3.2.5.4 ASM_POWER_LIMIT

The `ASM_POWER_LIMIT` initialization parameter specifies the default power for disk rebalancing in a disk group.

The range of values is 0 to 1024. The default value is 1. 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 and more rebalancing processes.

- For disk groups that have the disk group ASM compatibility set to 11.2.0.2 or higher (for example, `COMPATIBLE.ASM = 11.2.0.2`), the operational range of values is 0 to 1024 for the rebalance power.
- For disk groups that have the disk group ASM compatibility set to less than 11.2.0.2, the operational range of values is 0 to 11 inclusive. If the value for `ASM_POWER_LIMIT` is larger than 11, a value of 11 is used for these disk groups.

You can also specify the power of the rebalancing operation in a disk group with the `POWER` clause of the SQL `ALTER DISKGROUP .. REBALANCE` statement. The range of allowable values for the `POWER` clause is the same for the `ASM_POWER_LIMIT` initialization parameter. If the value of the `POWER` clause is specified larger than 11 for

a disk group with ASM compatibility set to less than 11.2.0.2, then a warning is displayed and a `POWER` value equal to 11 is used for rebalancing.

The specification of the power of the rebalancing operation in a disk group only affects rebalance operations, not new allocations to a disk group.

See Also:

- [Manually Rebalancing Disk Groups](#) (page 4-28) and [Tuning Rebalance Operations](#) (page 4-31) for information about the `ASM_POWER_LIMIT` initialization parameter, and the `POWER` clause
 - [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - *Oracle Database Reference* for more information about the `ASM_POWER_LIMIT` initialization parameter
 - *Oracle Database SQL Language Reference* for more information about the `POWER` clause of the `SQL ALTER DISKGROUP REBALANCE` statement
-
-

3.2.5.5 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.

Note:

The `ASM_PREFERRED_READ_FAILURE_GROUPS` initialization parameter has been deprecated in Oracle ASM 12c Release 2 (12.2) and may be desupported in a future release.

The `ASM_PREFERRED_READ_FAILURE_GROUPS` parameter setting is instance specific. The default value is `NULL`. This parameter is generally used for clustered Oracle ASM instances and its value can be different on different nodes.

For example:

```
diskgroup_name1.failure_group_name1,...
```

Note:

In an Oracle extended cluster, which contains nodes that span multiple physically separated sites, the `PREFERRED_READ.ENABLED` disk group attribute controls whether preferred read functionality is enabled for a disk group. If preferred read functionality is enabled, then this functionality enables an instance to determine and read from disks at the same site as itself, which can improve performance. Whether or not `PREFERRED_READ.ENABLED` has been enabled, preferred read can be set at the failure group level on an Oracle ASM instance or a client instance in a cluster with the `ASM_PREFERRED_READ_FAILURE_GROUPS` initialization parameter, which is available for backward compatibility.

The failure group values set by `ASM_PREFERRED_READ_FAILURE_GROUPS` take precedence over the `PREFERRED_READ.ENABLED` site values. The values of `ASM_PREFERRED_READ_FAILURE_GROUPS` on the client instances take precedence over the values on the Oracle ASM instances.

See Also:

- [Preferred Read Failure Groups](#) (page 4-46) for more information about preferred read failure groups
 - [PREFERRED_READ.ENABLED](#) (page 4-7) for information about the `PREFERRED_READ.ENABLED` disk group attribute
 - *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_GROUPS` initialization parameter
-

3.2.5.6 DB_CACHE_SIZE

The setting for the `DB_CACHE_SIZE` parameter determines the size of the buffer cache.

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

This buffer cache stores 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
-

3.2.5.7 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 for the Oracle Grid Infrastructure installation.

[Example 3-1](#) (page 3-14) shows an example of the diagnostic directory for an Oracle ASM instance.

Example 3-1 Sample diagnostic directory for an Oracle ASM instance

```
$ ls $ORACLE_BASE/diag/asm/+asm/+ASM
alert cdump hm incident incpkg ir lck metadata stage sweep trace
```

See Also:

- *Oracle Database Administrator's Guide* for more information about the `DIAGNOSTIC_DEST` initialization parameter and the fault diagnosability infrastructure.
 - *Oracle Database Reference* for more information about the `DIAGNOSTIC_DEST` parameter
-
-

3.2.5.8 INSTANCE_TYPE

The `INSTANCE_TYPE` initialization parameter specifies whether the instance is a database instance, an Oracle Automatic Storage Management (Oracle ASM) instance, or an Oracle ASM Proxy instance.

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

```
INSTANCE_TYPE = ASM
```

In addition to values `asm` and `rdbms`, `INSTANCE_TYPE` can be set to `asmproxy` in an Oracle Flex ASM configuration.

See Also:

- [Managing Oracle Flex ASM](#) (page 19-1) for information about Oracle Flex ASM
 - *Oracle Database Reference* for more information about the `INSTANCE_TYPE` parameter
-
-

3.2.5.9 LARGE_POOL_SIZE

The setting for the `LARGE_POOL_SIZE` parameter is used for large allocations.

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

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
-
-

3.2.5.10 PROCESSES

The `PROCESSES` initialization parameter affects Oracle ASM, but the default value is usually suitable.

However, if multiple database instances are connected to an Oracle ASM instance, then you can use the following formulas, where n is the number of database instances connecting to the Oracle ASM instance.

In a non-Exadata environment, the recommended settings are:

- For $n < 10$, `PROCESSES` = $50 * n + 50$
- For $n \geq 10$, `PROCESSES` = $10 * n + 450$

In an Oracle Exadata environment, the recommended setting is `PROCESSES` = `MAX(450 + 10 * n, 1024)`.

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
-
-

3.2.5.11 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
-
-

3.2.5.12 SHARED_POOL_SIZE

The setting for the `SHARED_POOL_SIZE` parameter determines the amount of memory required to manage the instance.

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

3.2.6 About 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.

The following are configuration 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 on Oracle ASM or 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
-
-

3.3 Administration of Oracle ASM Instances

Oracle ASM is typically installed in an Oracle Grid Infrastructure home separate from the Oracle Database home. Only one Oracle ASM instance is supported on a server in a standard configuration; however, Oracle Flex ASM provides additional configuration options.

When managing an Oracle ASM instance, the administration activity should be performed in the Oracle Grid Infrastructure home.

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

- [Using Oracle Grid Infrastructure for a Standalone Server](#) (page 3-17)
- [About Starting Up an Oracle ASM Instance](#) (page 3-18)
- [About Shutting Down an Oracle ASM Instance](#) (page 3-22)
- [About Administering Oracle ASM Instances with Server Control Utility](#) (page 3-23)
- [About Out of Place Upgrades](#) (page 3-23)
- [About Configuring Oracle Grid Infrastructure with the Configuration Wizard](#) (page 3-24)
- [About Active Session History Sampling for Oracle ASM](#) (page 3-24)
- [About Oracle Home User on Windows](#) (page 3-24)

See Also:

- [Managing Oracle Flex ASM](#) (page 19-1) for information about Oracle Flex ASM
 - [Creating Database Files in Oracle ASM Using SQL*Plus](#) (page 21-1) for information about Oracle ASM and creating database files
-
-

3.3.1 Using Oracle Grid Infrastructure for a Standalone Server

Oracle Grid Infrastructure for a standalone server (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 Grid Infrastructure for a standalone server (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 Database 2 Day + Real Application Clusters Guide* for information about automatically restarting single-instance databases residing on an Oracle RAC node
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for information about installation of Oracle Grid Infrastructure
-
-

3.3.2 About Starting Up an Oracle ASM Instance

This section describes how to start Oracle ASM instances under the following topics:

- [Connecting To and Starting Up an Oracle ASM Instance](#) (page 3-18)
- [Starting Up an Oracle ASM instance with an Incorrect SPFILE Path](#) (page 3-20)
- [About Mounting Disk Groups at Startup](#) (page 3-21)
- [About Restricted Mode](#) (page 3-21)

3.3.2.1 Connecting To and 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 with SQL*Plus, note the following:

- To connect to a local Oracle ASM instance with SQL*Plus, set the ORACLE_SID environment variable to the Oracle ASM system identifier (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 +ASM`node_number` where `node_number` is the number of the node. The ORACLE_HOME environment variable must be set to the Grid Infrastructure home where Oracle ASM was installed.

Note:

Oracle recommends that you do not change the default Oracle ASM SID name.

- 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.

After the Oracle ASM instance has started, you can mount disk groups with the `ALTER DISKGROUP . . . MOUNT` command.
- The associated Oracle Database instance does not have 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 encounter I/O errors and Oracle ACFS user data and metadata may not be written to storage before the Oracle ASM storage is fenced.
- **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. You can use the `RESTRICT` clause in combination with the `MOUNT`, `NOMOUNT`, and `OPEN` clauses.

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	

See Also:

- [Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#) (page 16-23) for information about dismounting an Oracle ACFS file system
 - [Authentication for Accessing Oracle ASM Instances](#) (page 3-27) for more information about user authentication
 - [startup](#) (page 10-29) for information about starting an Oracle ASM instance with `ASMCMD startup` command
 - [About Restricted Mode](#) (page 3-21) for information about restricted mode
 - [About Mounting Disk Groups at Startup](#) (page 3-21) for information about disk groups that are mounted at startup time
 - [Mounting and Dismounting Disk Groups](#) (page 4-67) for information about mounting disk groups
 - *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
-

3.3.2.2 Starting Up an Oracle ASM instance with an Incorrect SPFILE Path

If the SPFILE path in the GPnP profile is incorrect, you can start the Oracle ASM instance as follows:

1. Create a PFILE with one line in it that identifies the path to the SPFILE.

For example:

Create the `/oracle/dbs/spfileasm_init.ora` file that contains:

```
SPFILE= '+DATA/asm/asmparameterfile/asmspfile.ora'
```

2. Start up the instance using the initialization parameter file.

For example:

```
SQL> STARTUP PFILE=/oracle/dbs/spfileasm_init.ora
```

3. After the instance is running, use the ASMCMD `spset` command to update the SPFILE path in the GPnP profile. See "[spset](#) (page 10-28)".

For example:

```
ASMCMD> spset +DATA/asm/asmparameterfile/asmspfile.ora
```

See Also:

Oracle Database Administrator's Guide for more information about using STARTUP with a nondefault server parameter file

3.3.2.3 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](#) (page 4-67)".

3.3.2.4 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.

3.3.3 About Shutting Down an Oracle ASM Instance

You can initiate the Oracle ASM shutdown process by running 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 to the Grid Infrastructure home where Oracle ASM was installed.

If you are not using Oracle Flex ASM, 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. If you are using Oracle Flex ASM, Oracle Flex ASM clients move to other running Oracle ASM instances if an Oracle ASM instance is shut down.

If Oracle Cluster Registry (OCR) or voting files are stored in a disk group, the disk group can only be dismounted by shutting down the Oracle ASM instance as part of shutting down the clusterware on a node. To shut down the clusterware, run `crsctl stop crs`.

To shut down an Oracle ASM instance, perform the following steps:

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

The following list describes the `SHUTDOWN` modes and 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 behaves the same as `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 encounter I/O errors and Oracle ACFS user data and metadata may not be written to storage before the Oracle ASM storage is fenced.

See Also:

- [Authentication for Accessing Oracle ASM Instances](#) (page 3-27) for more information about user authentication
 - [Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#) (page 16-23) for information about dismounting an Oracle ACFS file system
 - [Authentication for Accessing Oracle ASM Instances](#) (page 3-27) for more information about user authentication on Oracle ASM instance
 - *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 managing voting files, Oracle Cluster Registry (OCR), and Oracle Local Registry (OLR)
-
-

3.3.4 About 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 the Oracle ASM Oracle Clusterware (CRS) resource in Oracle Cluster Registry (OCR)
 - Enable, disable, start, and stop Oracle ASM instances
 - 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

3.3.5 About Out of Place Upgrades

With an out-of-place upgrade, the installer installs the newer version of Oracle Grid Infrastructure in a separate Oracle Grid Infrastructure home.

An in-place upgrade of Oracle Grid Infrastructure 11g Release 2 (11.2) is not supported. For example, an upgrade of Oracle Grid Infrastructure 11g Release 2 (11.2.0.1) to Oracle Grid Infrastructure 11g Release 2 (11.2.0.2) must be an out of place upgrade.

See Also:

Oracle Grid Infrastructure Installation and Upgrade Guide for information about installing Oracle Grid Infrastructure, out of place upgrades, and performing rolling upgrades of Oracle Grid Infrastructure and Oracle ASM

3.3.6 About Configuring Oracle Grid Infrastructure with the Configuration Wizard

The Oracle Grid Infrastructure configuration wizard can update the configuration of an Oracle Grid Infrastructure environment after the software has been installed. The configuration wizard accepts your input, validates the input, and populates the configuration data into the `CRSCONFIG_PARAMS` file. If additional scripts must be run, the configuration wizard directs you to run those scripts.

See Also:

Oracle Clusterware Administration and Deployment Guide for information about the Oracle Grid Infrastructure configuration wizard.

3.3.7 About 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
-

3.3.8 About Oracle Home User on Windows

Oracle Database supports the use of an Oracle home user, which can be specified at installation time. The Oracle home user is associated with an Oracle home and it cannot be changed after installation. Different Oracle homes on a system can share the same Oracle home user or use different Oracle home user names.

In previous releases on Windows operating systems, Oracle services were required to run as `Local System` privileges, which are fully privileged. This feature enables the database, listener, and job scheduler services to run with low and non-administrative user privileges to allow tighter control of security. The Oracle home user can be a built-in account or a Windows user account. A Windows user account should be a low privileged (non-Administrator) account to ensure that the Oracle home user has a

limited set of privileges, ensuring that Oracle Database services have only those privileges required to run Oracle products. The Windows user account can be a Local User, a Domain User, or a Managed Services Account in general. However, Oracle RAC, Oracle Restart, and Oracle Grid Infrastructure installations require the use of the Domain User as the Oracle home user because a clusterwide identity is necessary.

See Also:

Oracle Database Platform Guide for Microsoft Windows for information about running Oracle services on Windows platforms and different types of Windows user accounts

3.4 Upgrading and Patching Oracle ASM

This section contains the following topics:

- [About Oracle ASM Rolling Upgrade](#) (page 3-25)
- [About Oracle ASM Rolling Patches](#) (page 3-26)

Note:

- For Oracle RAC environments, the Oracle Clusterware version number must be at least equal to the version number of the patch that you are applying to the Oracle Database.
 - You must apply the patch to the Oracle Grid Infrastructure home before you apply it to the Oracle Database home.
-

3.4.1 About Oracle ASM Rolling Upgrade

Oracle ASM rolling upgrade enables you to independently upgrade or patch clustered Oracle ASM nodes without affecting database availability which provides 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 an 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 Oracle ASM is upgraded with Oracle Clusterware for Oracle 11g Release 2 (11.2) or later as both are in the Oracle Grid Infrastructure home.

Note:

- Rolling upgrades only apply to clustered Oracle ASM instances, and you can only perform rolling upgrades on environments with Oracle Database 11g or later. 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 and Upgrade Guide* for information about performing, and restrictions related to, a rolling upgrade of Oracle ASM
 - *Oracle Database Upgrade Guide* for information about upgrading Oracle Database
 - *Oracle Database SQL Language Reference* for information about the rolling migration clause of the ALTER SYSTEM command
-
-

3.4.2 About Oracle ASM Rolling Patches

You can apply patches in a clustered Oracle ASM environment to update one node at a time to the latest patch level without affecting the overall availability of the Oracle ASM cluster or the database clusters using Oracle ASM for storage.

The ALTER SYSTEM ROLLING PATCH SQL statement enables you to start and stop rolling patches. For example:

```
SQL> ALTER SYSTEM START ROLLING PATCH;
```

```
SQL> ALTER SYSTEM STOP ROLLING PATCH;
```

You can determine if the cluster is in rolling patch mode by executing a SYS_CONTEXT SQL query for Cluster State. A new state (In Rolling Patch) is added to informing the user that the cluster is in rolling patch mode.

The queries in [Example 3-2](#) (page 3-27) display information about rolling patches. To run these queries, you must be connected to the Oracle ASM instance in the Grid home, and the Grid Infrastructure home must be configured with the Oracle Clusterware option for an Oracle RAC environment.

You can view all the patch Ids applied on the node and cluster by querying the V \$PATCHES view.

ASMCMD commands for rolling patches include:

- showclusterstate
- showpatches
- showversion

Example 3-2 Determining rolling patch mode and patch level

```
SELECT SYS_CONTEXT('SYS_CLUSTER_PROPERTIES', 'CLUSTER_STATE') FROM DUAL;  
  
SELECT SYS_CONTEXT('SYS_CLUSTER_PROPERTIES', 'CURRENT_PATCHLVL') FROM DUAL;
```

See Also:

- [ASMCMD Instance Management Commands](#) (page 10-11) for information about ASMCMD commands to monitor upgrade operations on an Oracle ASM instance
 - *Oracle Database SQL Language Reference* for information about the rolling patch clause of the ALTER SYSTEM command
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for information about performing, and restrictions related to, a rolling migration of Oracle ASM and Oracle Grid Infrastructure homes
-

3.5 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

This section describes the following topics:

- [About Privileges for Oracle ASM](#) (page 3-28)
- [Creating Users with the SYSASM Privilege](#) (page 3-31)
- [Operating System Authentication for Oracle ASM](#) (page 3-32)
- [Password File Authentication for Oracle ASM](#) (page 3-32)
- [Managing a Shared Password File in a Disk Group](#) (page 3-33)

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.

See Also:

- [Managing Oracle ASM File Access Control for Disk Groups](#) (page 4-61) for information about file permissions and Oracle ASM File Access Control
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for your operating system-specific information about how to ensure that the Oracle ASM and database instances have member disk access
 - *Oracle Database Security Guide* for information about maintaining database security, including assigning passwords
-

3.5.1 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 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](#) (page 3-28)
- [Using Separate Operating System Groups for Oracle ASM Users](#) (page 3-29)
- [The SYSASM Privilege for Administering Oracle ASM](#) (page 3-30)
- [The SYSDBA Privilege for Managing Oracle ASM Components](#) (page 3-31)

See Also:

[Oracle ACFS and File Access and Administration Security](#) (page 11-16) for information about privileges and Oracle ACFS

3.5.1.1 Using One Operating System Group for Oracle ASM Users

If you do not want to divide the privileges for system 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](#) (page 3-29) 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

3.5.1.2 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** (for example: `asmadmin`)
This group is granted the `SYSASM` privilege, which provides full administrative privileges for the Oracle ASM instance.
- **OSDBA for Oracle ASM group** (for example: `asmdba`)
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 (`dba`).
- **OSOPER for Oracle ASM group** (for example: `asmoper`)
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.

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 Grid Infrastructure, the Oracle ASM user, such as `grid`, does not have 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`) must 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 the owner of the Oracle ASM software.

[Table 3-2](#) (page 3-30) 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)

3.5.1.3 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 SYSASM privilege with password file authentication, as discussed in "[Password File Authentication for Oracle ASM](#) (page 3-32)".

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
```

3.5.1.4 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 a 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

3.5.2 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 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.
 - Oracle ASM does not support the creation of external and global users.
-
-

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;
```

When you revoke the last privilege of a user in an Oracle ASM password file, the user is not automatically deleted as is done in the Oracle Database password file. You must run `DROP USER` to delete a user with no privileges in an Oracle ASM password file.

See Also:

- [orapwusr](#) (page 10-16) for information about creating a user with Oracle ASM command-line utility (`ASMCMD`)
 - *Oracle Database SQL Language Reference* for information about the `CREATE USER` and `GRANT` commands
 - *Oracle Database Administrator's Guide* for information about creating and maintaining a password file
-

3.5.3 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, `dba` is the default operating system group designated as OSASM, OSOPER, and OSDBA for Oracle ASM.

On Windows systems, `ORA_ASMADMIN`, `ORA_ASMDBA`, and `ORA_ASMOPER` are the operating system groups designated for OSASM, OSDBA and OSOPER respectively for Oracle ASM.

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 and Upgrade Guide* for information about installation of the Oracle Grid Infrastructure
 - *Oracle Database Installation Guide* for your specific operating system for information about default operating system groups
-

3.5.4 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.

If you select the Oracle ASM storage option, then `ASMCA` creates a password file for Oracle ASM with initial users (`SYS` and `ASMSNMP`) when `ASMCA` configures the

Oracle ASM disk groups. To add other users to the password file, you can use the `CREATE USER` and `GRANT` commands.

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 use password file authentication.

See Also:

- [About Privileges for Oracle ASM](#) (page 3-28) for information about privileges for Oracle ASM
 - *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$PWFILERS` view which lists users who have been granted `SYSASM`, `SYSDBA`, and `SYSOPER` privileges as derived from the password file.
-
-

3.5.5 Managing a Shared Password File in a Disk Group

You can manage a shared password file on a disk group enabling access to the single file across the cluster system.

This sections contains these topics:

- [About a Shared Password File in a Disk Group](#) (page 3-34)
- [Creating a Password File in a Disk Group](#) (page 3-34)
- [Backing Up and Restoring an Oracle ASM Password File in a Disk Group](#) (page 3-35)

For information about `ASMCMD` commands to manage an Oracle ASM or database instance password file in a disk group; such as `pwcopy`, `pwcreate`, and `pwmove`; refer to "[ASMCMD Instance Management Commands](#) (page 10-11)".

See Also:

- *Oracle Database Administrator's Guide* for information about maintaining a password file
 - *Oracle Database Reference* for information about the `REMOTE_LOGIN_PASSWORDFILE` initialization parameter
 - *Oracle Real Application Clusters Administration and Deployment Guide* for information about `SRVCTL`
-
-

3.5.5.1 About a Shared Password File in a Disk Group

An individual password file for Oracle Database or Oracle ASM can reside on a designated Oracle ASM disk group. Having the password files reside on a single location accessible across the cluster reduces maintenance costs and situations where passwords become out of sync.

You can use a password file located on a disk group for authentication only if the Oracle ASM instance is running and the designated disk group is mounted. Otherwise, operating system authentication must be used to bootstrap the startup of the Oracle ASM instance and stack.

The `COMPATIBLE .ASM` disk group attribute must be set to at least 12.1 for the disk group where the password is to be located. The `SYSASM` privilege is required to manage the Oracle ASM password file. The `SYSDBA` privilege on Oracle ASM is required to manage the database password file.

The shared password file in a disk group is managed by `ASMCMD` commands, the `ORAPWD` tool, and `SRVCTL` commands. `ORAPWD` supports the creation of password files on an Oracle ASM disk group. All other password file manipulation is performed with `ASMCMD` or `SRVCTL` commands.

Before running commands, such as `ORAPWD`, to create a password file, ensure that the `ORACLE_SID` and `ORACLE_HOME` environmental variables have been set properly. For example, before setting the password file for Oracle ASM, set the `ORACLE_SID` and `ORACLE_HOME` environmental variables to ensure that you can connect to the local Oracle ASM instance. For information about environmental variables and connecting to an Oracle ASM instance, refer to "[Connecting To and Starting Up an Oracle ASM Instance](#) (page 3-18)".

`SRVCTL` provides commands to manage a password file in a disk group, such as the following commands for updating and displaying the location of the password file:

```
$ srvctl modify asm -pwfile location
$ srvctl modify database -db dbname -pwfile location
$ srvctl config asm
```

3.5.5.2 Creating a Password File in a Disk Group

Note:

Do not use `orapwd` to create a new Oracle ASM password file when Oracle ASM is configured in a cluster. If the Oracle ASM password file is inaccessible in a cluster configuration, you must restore the password file using an available backup rather than create a new file. For articles about recreating a shared Oracle ASM password file in a cluster, such as document 1929673.1, refer to articles at [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com).

When using `orapwd` to create a database password file in a disk group, you must specify the disk group location and database unique name.

For example:

```
$ orapwd file='+data/ORCL/orapwdb' dbuniquename='orcl'
```

Enter password for SYS:

The `asm` switch specifies that `orapwd` create an Oracle ASM password file rather than a database password file.

For example:

```
$ orapwd file='+data/ASM/orapwasm' asm=y
```

Enter password for SYS:

You can create a new password file in a disk group using a password file from a previous release.

For example:

```
$ orapwd input_file='/oraclegrid/dbs/orapwasm' file='+data/ASM/orapwasm' asm=y
```

Enter password for SYS:

3.5.5.3 Backing Up and Restoring an Oracle ASM Password File in a Disk Group

You can make a backup of the Oracle ASM password file, and if the Oracle ASM password file is lost or the disk group becomes inaccessible, then you can restore the backup password file.

This section describes the steps to back up the Oracle ASM password file to a disk group and the steps to restore the Oracle ASM password file.

The source and target disk groups must have the disk group attribute `COMPATIBLE.ASM` set to 12.1 or higher.

1. Locate the password file using the `ASMCMD pwget` command.

For example:

```
ASMCMD [+] > pwget --asm
+DATA/orapwasm
```

2. Back up the password file to another disk group with the `pwcopy` command.

For example:

```
ASMCMD [+] > pwcopy +DATA/orapwasm +FRA/my_pwfile_backup
```

Using `pwcopy` without the `--asm` or `--dbuniquefilename` option does not change the current location of the password file. If necessary after the copy is made, you can set the current password file location with the `pwset` command.

3. Verify which password file is in the current location after making a backup with the `pwcopy` command.

For example:

```
ASMCMD [+] > pwget --asm
+DATA/orapwasm
```

4. Verify the backup password file was created.

For example:

```
ASMCMD [+] > ls +fra/my_pwfile_backup
my_pwfile_backup
```

5. To restore the Oracle ASM password file, you can use `pwset` or `pwcopy`.

To restore the Oracle ASM password file from the backup and use the existing location, use the `pwset` command with the `--asm` option.

For example:

```
ASMCMD [+] > pwset --asm +FRA/my_pwfile_backup
```

To restore the Oracle ASM password file from the backup to another disk group, use the `pwcopy` command with the `--asm` option.

For example:

```
ASMCMD [+] > pwcopy --asm +FRA/my_pwfile_backup +DATA2/my_orapwasm
```

The `--asm` option with the `pwset` and `pwcopy` command sets the password location for the Oracle ASM instance.

6. Verify the location of the current password file with the `pwget` command if you have changed the location.

For example:

```
ASMCMD [+] > pwget --asm  
+DATA2/my_orapwasm
```

3.6 Administering the Oracle ASM Audit Trail

The number of audit trail files in the audit destination directories for an Oracle ASM instance can grow very large if the directories are not regularly maintained. To control the number of these files, Oracle ASM auditing can be managed with operating system tools, such as the Syslog facility on UNIX platforms.

Managing Oracle ASM Audit Records With Syslog

The Oracle ASM audit records directed to the Syslog facility should remain separated from other system generated audit records in the system. To ensure that separation, set the configuration values in `/etc/syslog.conf` so that only Oracle audit records are written to a given file.

For example, you could choose to set the `/var/log/oracle/oracleaudit.log` file exclusively for Oracle audit records with the following setting in the `syslog.conf` file:

```
# Log all Oracle audit records.  
LOCAL7.ALERT    /var/log/oracle/oracleaudit.log
```

The syslog daemon should be restarted to pick up the changes in the syslog configuration file. The restart operation requires super user (`root`) privileges on the machine. For example:

```
# /etc/rc.d/init.d/syslog restart
```

After setting up the entry in the syslog configuration file, set the `AUDIT_SYSLOG_LEVEL` initialization parameter in the Oracle ASM instance parameter file to the same value (`AUDIT_SYSLOG_LEVEL = LOCAL7.ALERT`) and restart the Oracle ASM instance.

See Also:

- Articles at [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com) for information about managing Oracle ASM auditing. For example:
 - *Manage ASM Audit Files with syslog* (Doc ID 1559573.1)
 - *Manage Audit File Directory Growth with cron* (Doc ID 1298957.1)
 - *AUDIT_SYS_OPERATIONS Set To FALSE Yet Audit Files Are Generated* (308066.1)
 - *Init.ora Parameter "AUDIT_FILE_DEST" Reference Note* (39796.1)
 - *Oracle Database Reference* for information about the `AUDIT_SYSLOG_LEVEL` initialization parameter.
-

3.7 Migration of 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, 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. To migrate to Oracle ASM, you can use the methods described in the following sections:

- [About Using Oracle Recovery Manager to Migrate Databases to Oracle ASM](#) (page 3-37)
- [Best Practices White Papers on Migrating to Oracle ASM](#) (page 3-37)

Note:

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

3.7.1 About 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 [Performing Oracle ASM Data Migration with RMAN](#) (page 8-1).

3.7.2 Best Practices White Papers on Migrating to Oracle ASM

The Oracle Maximum Availability Architecture (MAA) website 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 documentation at the MAA link on Oracle Technology Network:

<http://www.oracle.com/technetwork/database/features/availability/maa-096107.html>

Administering Oracle ASM Disk Groups

Administration of Oracle ASM disk groups includes creating, altering, dropping, mounting, and dismounting tasks.

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:

- [Managing Disk Group Attributes](#) (page 4-2)
- [Creating Disk Groups](#) (page 4-9)
- [Altering Disk Groups](#) (page 4-19)
- [Oracle ASM Disk Discovery](#) (page 4-33)
- [Managing Capacity in Disk Groups](#) (page 4-35)
- [Oracle ASM Mirroring and Disk Group Redundancy](#) (page 4-38)
- [Performance and Scalability Considerations for Disk Groups](#) (page 4-48)
- [Disk Group Compatibility](#) (page 4-50)
- [Managing Oracle ASM File Access Control for Disk Groups](#) (page 4-61)
- [Maintaining Disk Groups](#) (page 4-66)
- [Managing Oracle ASM Flex Disk Groups](#) (page 4-72)

See Also:

- [About Starting Up an Oracle ASM Instance](#) (page 3-18) for information about starting up an Oracle ASM instance
 - [Managing Oracle ASM With ASMCA](#) (page 9-1) for information about administering Oracle ASM disk groups with Oracle ASM Configuration Assistant (ASMCA)
 - [Managing Oracle ASM with ASMCMD](#) (page 10-1) for information about administering Oracle ASM disk groups with ASMCMD
 - The Oracle Cloud Storage page on the Oracle Technology Network website at <http://www.oracle.com/technetwork/database/cloud-storage/index.html> for more information about Oracle ASM
-

4.1 Managing Disk Group Attributes

Disk group attributes are parameters that are bound to a disk group, rather than an Oracle ASM instance. Some disk group attributes can be set when a disk group is created or altered. Other disk attributes can only be set when a disk group is created or only be set when a disk group is altered.

This section contains the following topics:

- [About Viewing and Setting Disk Group Attributes](#) (page 4-2)
- [Summary of Disk Group Attributes](#) (page 4-3)

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](#) (page 5-19)".

4.1.1 About Viewing and Setting Disk Group Attributes

You can set and display disk group attribute settings.

You can display disk group attributes with the `V$ASM_ATTRIBUTE` view and the `ASMCMD lsattr` command.

Disk group attributes can be set with the `ATTRIBUTE` clause of the `ALTER DISKGROUP` or `CREATE DISKGROUP SQL` statement and the `ASMCMD setattr` and `mkdg` commands.

See Also:

- [Example 6-1](#) (page 6-3) for an example of the use of the `V $ASM_ATTRIBUTE` view
 - [lsattr](#) (page 10-60) for information about the `lsattr` command
 - [setattr](#) (page 10-75) for information about the `setattr` command
 - [mkdgm](#) (page 10-69) for information about the `mkdgm` command
 - [Managing Disk Groups with ASMCA](#) (page 9-2) for information about setting disk group attributes with Oracle ASM Configuration Assistant (ASMCA)
-

4.1.2 Summary of Disk Group Attributes

The Oracle ASM disk group attributes are listed in this section.

- [ACCESS_CONTROL.ENABLED](#) (page 4-4)
- [ACCESS_CONTROL.UMASK](#) (page 4-4)
- [AU_SIZE](#) (page 4-4)
- [CELL.SMART_SCAN_CAPABLE](#) (page 4-4)
- [CELL.SPARSE_DG](#) (page 4-4)
- [COMPATIBLE.ASM](#) (page 4-4)
- [COMPATIBLE.RDBMS](#) (page 4-4)
- [COMPATIBLE.ADVANCED](#) (page 4-5)
- [CONTENT.CHECK](#) (page 4-5)
- [CONTENT.TYPE](#) (page 4-5)
- [DISK_REPAIR_TIME](#) (page 4-6)
- [FAILGROUP_REPAIR_TIME](#) (page 4-6)
- [IDP.BOUNDARY](#) and [IDP.TYPE](#) (page 4-6)
- [LOGICAL_SECTOR_SIZE](#) (page 4-7)
- [PHYS_META_REPLICATED](#) (page 4-7)
- [PREFERRED_READ.ENABLED](#) (page 4-7)
- [SCRUB_ASYNC_LIMIT](#) (page 4-8)
- [SECTOR_SIZE](#) (page 4-8)
- [STORAGE.TYPE](#) (page 4-8)
- [THIN_PROVISIONED](#) (page 4-9)

4.1.2.1 ACCESS_CONTROL.ENABLED

This attribute configures Oracle ASM File Access Control and can only be set when altering a disk group.

For information about the `ACCESS_CONTROL.ENABLED` attribute, see "[Setting Disk Group Attributes for Oracle ASM File Access Control](#) (page 4-63)".

4.1.2.2 ACCESS_CONTROL.UMASK

This attribute configures Oracle ASM File Access Control and can only be set when altering a disk group.

For information about the `ACCESS_CONTROL.UMASK` attribute, see "[Setting Disk Group Attributes for Oracle ASM File Access Control](#) (page 4-63)".

4.1.2.3 AU_SIZE

This attribute specifies allocation unit (AU) size and can only be set when creating a disk group. [Example 4-1](#) (page 4-13) shows how the `AU_SIZE` is specified with the `CREATE DISKGROUP` SQL statement.

For information about allocation units, refer to "[About Oracle ASM Allocation Units](#) (page 1-7)".

4.1.2.4 CELL.SMART_SCAN_CAPABLE

This attribute is only applicable to the configuration of Oracle Exadata storage.

See Also:

Oracle Exadata documentation

4.1.2.5 CELL.SPARSE_DG

This attribute is only applicable to the configuration of Oracle Exadata storage.

See Also:

Oracle Exadata documentation

4.1.2.6 COMPATIBLE.ASM

This attribute specifies the Oracle ASM compatibility level of a disk group. Compatibility attributes can be set when the disk group is created and advanced when a disk group is altered, but compatibility attributes cannot be reversed. For information about the `COMPATIBLE.ASM` attribute, see "[COMPATIBLE.ASM](#) (page 4-52)".

4.1.2.7 COMPATIBLE.RDBMS

This attribute specifies the database compatibility level of a disk group. Compatibility attributes can be set when the disk group is created and advanced when a disk group is altered, but compatibility attributes cannot be reversed. For information about the `COMPATIBLE.RDBMS` attribute, see "[COMPATIBLE.RDBMS](#) (page 4-52)".

4.1.2.8 COMPATIBLE.ADVM

This attribute specifies the Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) compatibility level of a disk group. Compatibility attributes can be set when the disk group is created and advanced when a disk group is altered, but compatibility attributes cannot be reversed. For information about the `COMPATIBLE.ADVM` attribute, see "[COMPATIBLE.ADVM](#) (page 4-52)" and "[Overview of Oracle ASM Dynamic Volume Manager](#) (page 11-49)".

4.1.2.9 CONTENT.CHECK

The `CONTENT.CHECK` attribute enables or disables content checking when performing data copy operations for rebalancing a disk group.

The attribute value can be set to `true` or `false`.

The content checking can include Hardware Assisted Resilient Data (HARD) checks on user data, validation of file types from the file directory against the block contents and file directory information, and mirror side comparison.

When the attribute is set to `true`, logical content checking is enabled for all rebalance operations.

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

See Also:

[Manually Rebalancing Disk Groups](#) (page 4-28) and [Tuning Rebalance Operations](#) (page 4-31) for information about disk group rebalancing

4.1.2.10 CONTENT.TYPE

This attribute identifies the disk group type: `data`, `recovery`, or `system`. The type value determines where Oracle ASM places the nearest neighbor disk in the failure group where Oracle ASM mirrors copies of the data. Oracle ASM uses this attribute value to make it less likely that a double failure in the storage medium causes disk groups of different content types to become unavailable. Specifying this attribute enables Oracle ASM to better deploy all disk groups in the available storage.

The default value is `data` which specifies a distance of 1 to the nearest neighbor disk. A value of `recovery` specifies a distance of 3 to the nearest neighbor disk and a value of `system` specifies a distance of 5.

The attribute can be specified when creating or altering a disk group. If `CONTENT.TYPE` attribute is set or changed using `ALTER DISKGROUP`, then the new configuration does not take effect until a disk group rebalance is explicitly run.

The `CONTENT.TYPE` attribute is only valid for disk groups that are set to normal or high redundancy. The `COMPATIBLE.ASM` attribute must be set to `11.2.0.3` or higher to enable the `CONTENT.TYPE` attribute for the disk group.

This attribute can be set with `ASMCMD mkgg` or `SQL CREATE` and `ALTER DISKGROUP` statements.

For example:

```
SQL> CREATE DISKGROUP data NORMAL REDUNDANCY
      FAILGROUP controller1 DISK
```

```

'/devices/disk1' NAME disk1,
...
ATTRIBUTE
'compatible.asm' = '12.1.0.1',
'content.type' = 'recovery',
...

SQL> ALTER DISKGROUP data SET ATTRIBUTE 'content.type' = 'data';

```

This attribute is primarily intended for use with Oracle Exadata storage.

For information about disk resync, see "[Oracle ASM Fast Mirror Resync](#) (page 4-43)".

See Also:

- *Oracle Exadata documentation*
 - [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com) for information about using the CONTENT.TYPE attribute with Oracle Exadata storage
-
-

4.1.2.11 DISK_REPAIR_TIME

This attribute specifies the time interval to repair a disk and bring it back online before initiating the drop operation. This attribute can only be set when altering a disk group and is only applicable to normal and high redundancy disk groups.

For information about disk repair time, see "[Oracle ASM Fast Mirror Resync](#) (page 4-43)".

4.1.2.12 FAILGROUP_REPAIR_TIME

This attribute specifies a default repair time for the failure groups in the disk group. The failure group repair time is used if Oracle ASM determines that an entire failure group has failed. The default value is 24 hours (24h). If there is a repair time specified for a disk, such as with the DROP AFTER clause of the SQL ALTER DISKGROUP OFFLINE DISK statement, that disk repair time overrides the failure group repair time.

This attribute can only be set when altering a disk group and is only applicable to normal and high redundancy disk groups.

For information about failure group repair time, see "[Oracle ASM Fast Mirror Resync](#) (page 4-43)".

4.1.2.13 IDP.BOUNDARY and IDP.TYPE

Note:

The Intelligent Data Placement (IDP) feature has been deprecated in Oracle ASM 12c Release 2 (12.2) and may be desupported in a future release.

These attributes are used to configure Oracle Exadata storage.

See Also:

Oracle Exadata documentation

4.1.2.14 LOGICAL_SECTOR_SIZE

The `LOGICAL_SECTOR_SIZE` disk group attribute specifies the logical sector size for a disk group. The value can be set when creating or altering a disk group.

The value of `LOGICAL_SECTOR_SIZE` specifies the smallest unit of I/O allowed on the disk group. The value must be less than or equal to the size of the `SECTOR_SIZE` disk group attribute.

`COMPATIBLE.ASM` must be set to 12.2 or greater to use the `LOGICAL_SECTOR_SIZE` disk group attribute.

For information about setting the `LOGICAL_SECTOR_SIZE` disk group attribute, see "[Specifying the Logical and Physical Sector Size](#) (page 4-15)".

For information about the `SECTOR_SIZE` disk group attribute, refer to [SECTOR_SIZE](#) (page 4-8).

4.1.2.15 PHYS_META_REPLICATED

This disk group attribute tracks the replication status of a disk group. When the Oracle ASM compatibility of a disk group is advanced to 12.1 or higher, the physical metadata of each disk is replicated. This metadata includes the disk header, free space table blocks, and allocation table blocks. The replication is performed online asynchronously. This attribute value is set to `true` by Oracle ASM if the physical metadata of every disk in the disk group has been replicated.

This disk group attribute is only defined in a disk group with the Oracle ASM disk group compatibility (`COMPATIBLE.ASM`) set to 12.1 and higher. This attribute is read-only and is intended for information only. A user cannot set or change its value. The values are either `true` or `false`.

4.1.2.16 PREFERRED_READ_ENABLED

In an Oracle extended cluster, which contains nodes that span multiple physically separated sites, the `PREFERRED_READ.ENABLED` disk group attribute controls whether preferred read functionality is enabled for a disk group.

If preferred read functionality is enabled, then this functionality enables an instance to determine and read from disks at the same site as itself, which can improve performance.

Each Oracle ASM instance in an Oracle extended cluster has the information to determine the site in which it is located. If preferred read functionality is enabled, then based on the site information, the preferred status for a disk is determined to be `true` if the disk is in the same site as the instance. Otherwise, if the disk is not located at the same site as the Oracle ASM instance, the preferred status for a disk is set to `false`.

For extended clusters, the preferred read setting is enabled by default with a `TRUE` setting for the `PREFERRED_READ.ENABLED` disk group attribute of each disk group. For clusters that are not extended (only one physical site), preferred read is disabled. Preferred read status applies to extended, normal, high, and flex redundancy disk groups.

Oracle Database 12c Release 2 (12.2) is required and `COMPATIBLE .ASM` must be set to 12 . 2 or greater to use `PREFERRED_READ . ENABLED`.

Whether or not `PREFERRED_READ . ENABLED` has been enabled, preferred read can be set at the failure group level on an Oracle ASM instance or a client instance in a cluster with the `ASM_PREFERRED_READ_FAILURE_GROUPS` initialization parameter, which is available for backward compatibility.

To determine whether a disk has preferred read status, you can check the `PREFERRED_READ` column of the `V$ASM_DISK` view.

See Also:

- [ASM_PREFERRED_READ_FAILURE_GROUPS](#) (page 3-12) for information about the `ASM_PREFERRED_READ_FAILURE_GROUPS` initialization parameter
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for information about Oracle extended clusters
 - *Oracle Database Reference* for information about all of the `V$ASM*` dynamic performance views
-
-

4.1.2.17 SCRUB_ASYNC_LIMIT

The `SCRUB_ASYNC_LIMIT` disk group attribute sets the maximum allowed outstanding parallel read requests for asynchronous scrubbing.

The range of values is 1 to 1024. The default value is 1.

`COMPATIBLE .ASM` must be set to 12 . 2 or greater to use this disk group attribute.

For information about the scrubbing operation, refer to [Scrubbing Disk Groups](#) (page 4-32).

4.1.2.18 SECTOR_SIZE

The `SECTOR_SIZE` disk group attribute specifies the sector size for a disk group. The value can be set when creating or altering a disk group.

The value of `SECTOR_SIZE` is used for the block size of newly-created redo logs and Oracle ACFS file systems. The value can be changed to facilitate migration of disks to different sector sizes. The value must be greater than or equal to the size of the `LOGICAL_SECTOR_SIZE` disk group attribute.

`COMPATIBLE .ASM` must be set to 12 . 2 or greater to use the 12.2 functionality of `SECTOR_SIZE`, which includes the ability to alter the sector size after a disk group has been created.

For information about setting the `SECTOR_SIZE` disk group attribute, see "[Specifying the Logical and Physical Sector Size](#) (page 4-15)".

For information about the `LOGICAL_SECTOR_SIZE` disk group attribute, refer to [LOGICAL_SECTOR_SIZE](#) (page 4-7).

4.1.2.19 STORAGE.TYPE

This disk group attribute specifies the type of the disks in the disk group. The possible values are `AXIOM`, `ZFSSA`, and `OTHER`. If the attribute is set to `AXIOM` or `ZFSSA`, then

all disks in the disk group must be of that type. If the attribute is set to `OTHER`, then any types of disks can be in the disk group.

If the `STORAGE . TYPE` disk group attribute is set to `AXIOM` or `ZFSSA`, then functionality for Hybrid Columnar Compression (HCC) can be enabled for Pillar Axiom or ZFS storage.

Note:

Pillar Axiom storage must be added as SCSI/Fiber Channel.

To set the `STORAGE . TYPE` attribute, the `COMPATIBLE . ASM` and `COMPATIBLE . RDBMS` disk group attributes must be set to `11 . 2 . 0 . 3` or higher. For maximum support with ZFS storage, set the `COMPATIBLE . ASM` and `COMPATIBLE . RDBMS` disk group attributes to `11 . 2 . 0 . 4` or higher. For information about disk group compatibility attributes, refer to "[Disk Group Compatibility](#) (page 4-50)".

The `STORAGE . TYPE` attribute can be set when creating a disk group or when altering a disk group. The attribute cannot be set when clients are connected to the disk group. For example, the attribute cannot be set when an Oracle ADVM volume is enabled on the disk group.

The attribute is not visible in the `V$ASM_ATTRIBUTE` view or with the `ASMCMD lsattr` command until the attribute has been set.

See Also:

Oracle Database Concepts for more information about Hybrid Columnar Compression. Hybrid Columnar Compression is a feature of certain Oracle storage systems.

4.1.2.20 THIN_PROVISIONED

This attribute enables or disables the functionality to discard unused storage space after a disk group rebalance is completed. The attribute value can be `true` or `false`. The default value is `false`.

Storage vendor products that support thin provisioning have the capability to reuse the discarded storage space for a more efficient overall physical storage utilization.

Note:

The `THIN_PROVISIONED` attribute is supported only in Oracle Database 12c Release 1 with Oracle ASM Filter Driver (Oracle ASMFDF).

For information about disk group rebalancing, refer to "[Tuning Rebalance Operations](#) (page 4-31)".

4.2 Creating Disk Groups

This topic introduces 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](#) (page 4-10)
- [Creating Disk Groups for Data and Fast Recovery Area](#) (page 4-12)
- [Creating Disk Groups for a New Oracle Installation](#) (page 4-14)
- [Specifying the Logical and Physical Sector Size](#) (page 4-15)
- [Storing Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#) (page 4-18)

See Also:

- *Oracle Grid Infrastructure Installation and Upgrade Guide* for information about storage requirements for Oracle Automatic Storage Management
 - *Oracle Database SQL Language Reference* for information about the CREATE DISKGROUP SQL statement
-

4.2.1 Using the CREATE DISKGROUP SQL Statement

You can use CREATE DISKGROUP SQL statement 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 Oracle Database object names that must be quoted. While the use of quoted identifiers may be valid as names in a SQL statement, such as CREATE DISKGROUP "1DATA", the names may not be valid when using other tools that manage the object. For more information about naming Oracle Database objects, refer to *Oracle Database SQL Language Reference*.

- Specify the redundancy level of the disk group.

For Oracle ASM to mirror files, specify a redundancy level such as NORMAL REDUNDANCY (two-way mirroring by default for most file types), HIGH REDUNDANCY (three-way mirroring for all files), FLEX REDUNDANCY, or EXTENDED REDUNDANCY.

Specify EXTERNAL REDUNDANCY if you do not want mirroring by Oracle ASM. For example, you might choose EXTERNAL REDUNDANCY to use storage array protection features.

After a disk group is created with NORMAL or HIGH redundancy, you can migrate the disk group to a FLEX redundancy disk group with the CONVERT REDUNDANCY TO FLEX clause of the ALTER DISKGROUP SQL statement. You cannot migrate an EXTERNAL disk group to a FLEX redundancy disk group. When migrating to a flex disk group, the disk group must have a minimum of three failure groups and must be mounted in restricted mode (MOUNTED RESTRICTED) for the duration of the migration.

To change the redundancy level for NORMAL, HIGH, or EXTERNAL disk groups to a NORMAL, HIGH, or EXTERNAL, you must create another disk group with the appropriate redundancy and then move the files to the new disk group.

For FLEX and EXTENDED redundancy disk groups, the disk group redundancy cannot be changed. However, the redundancy of the files in the disk group can be changed by modifying the REDUNDANCY property of the respective file group.

- 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.

You must create failure groups of equal size to maintain space balance and even distribution of mirror data.

- Optionally specify the type of failure group, such as QUORUM or REGULAR.
- Optionally specify disk group attributes, such as software compatibility or allocation unit size.

All disks in the disk group must be of equal size. Oracle ASM programmatically determines the size of each disk. If for some reason the size cannot be determined, or to restrict the amount of space used on a disk, you can specify a SIZE clause for a disk; however, the same size value must be specified for every disk in the disk group.

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.

You can use the ASMCMD stamp command to provision disks for Oracle ASM before creating a disk group that uses them. If a disk is provisioned, then the disk name, failure group, and site information are specified at the time the disk is provisioned and do not need to be specified when the disk group is created.

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 with FORCE. If a disk is dropped with NOFORCE, then you can add it with 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. 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 you shut down and restart the Oracle ASM instance. You can also create disk groups with Oracle ASM Configuration Assistant (ASMCA) and the `ASMCMD mkg` command.

See Also:

- [Managing Oracle ASM Flex Disk Groups](#) (page 4-72) for information about flex disk groups and file groups
 - [Mirroring, Redundancy, and Failure Group Options](#) (page 4-38) and [Understanding Oracle ASM Concepts](#) (page 1-3) for information about redundancy levels and failure groups
 - [Storing Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#) (page 4-18) for information about `QUORUM` and `REGULAR` failure groups
 - `stamp` (page 10-76) for information about the `ASMCMD stamp` command
 - [Managing Disk Groups with ASMCA](#) (page 9-2) for information about creating disk groups with Oracle ASM Configuration Assistant (ASMCA)
 - `mkg` (page 10-69) for information about creating disk groups with the `ASMCMD mkg` command
 - *Oracle Database SQL Language Reference* for information about the `CREATE DISKGROUP` SQL statement
-

4.2.2 Creating Disk Groups for Data and Fast Recovery Area

You can create a disk group for data and another disk group for the fast recovery area.

The SQL statement in [Example 4-1](#) (page 4-13) 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.

The example assumes that the `ASM_DISKSTRING` initialization parameter is set to the `'/devices/*'` string. Oracle ASM disk discovery identifies `diska1` through `diska4` and `diskb1` through `diskb4` in the `/devices` directory.

In [Example 4-1](#) (page 4-13), 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.

When creating the disk group in [Example 4-1](#) (page 4-13), the values of following disk group attributes were explicitly set:

- `AU_SIZE`
Specifies the size of the allocation unit for the disk group.

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.
- `COMPATIBLE.RDBMS`
Determines the minimum software version for any database instance that uses a disk group.
- `COMPATIBLE.ADVM`
Determines whether the disk group can contain Oracle ADVM volumes.

In [Example 4-2](#) (page 4-13), 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 example assumes that `diskc1` through `diskc9` are present in the `/devices` directory.

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';
```

Example 4-2 Creating the FRA disk group

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

See Also:

- [Administering Oracle ASM Filter Driver](#) (page 20-1) for information about the disk name format when using Oracle ASM Filter Driver
 - [About Oracle ASM Allocation Units](#) (page 1-7) for information about allocation units
 - [COMPATIBLE.ASM](#) (page 4-52) for information about the COMPATIBLE.ASM attribute
 - [COMPATIBLE.RDBMS](#) (page 4-52) for information about the COMPATIBLE.RDBMS attribute
 - [COMPATIBLE.ADVM](#) (page 4-52) for information about the COMPATIBLE.ADVM attribute
-

4.2.3 Creating Disk Groups for a New Oracle Installation

This topic describes the basic procedure for creating disk groups during a new installation of Oracle Grid Infrastructure standalone (Oracle Restart) option and Oracle Database.

The procedure described in this topic for a new installation of Oracle Grid Infrastructure standalone (Oracle Restart) assumes that a disk group is configured for storing the database data files and a disk group is configured for storing the fast recovery area files.

Note:

In Oracle Grid Infrastructure for cluster configurations, the Oracle Cluster Registry (OCR) and voting files should be stored on their own disk group. However, if only two disk groups are created for the Oracle ASM installation, then store the OCR and voting files in the disk group that contains the database data files.

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

During the installation, create a disk group (`data`) for storing database data files and control files.

This OUI disk group creation process is similar to creating a disk group with Oracle ASM Configuration Assistant (ASMCA). 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.

2. After the Oracle Grid Infrastructure is installed, use ASMCA to create a disk group (`fra`) for storing the fast recovery area files.

You can also create the `fra` disk group with SQL*Plus or ASMCMD commands run from the Oracle Grid Infrastructure home.

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:

- [Managing Disk Groups with ASMCA](#) (page 9-2) for information about using ASMCA to create or alter disk groups
 - [Storing Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#) (page 4-18) for information about storing Oracle Cluster Registry (OCR) and voting files
 - [Using the CREATE DISKGROUP SQL Statement](#) (page 4-10) for information about creating a disk group with SQL*Plus
 - [mkdgm](#) (page 10-69) for information about creating a disk group with ASMCMD
 - *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
 - *Oracle Grid Infrastructure Installation and Upgrade 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
-
-

4.2.4 Specifying the Logical and Physical Sector Size

You can use the optional `LOGICAL_SECTOR_SIZE` and `SECTOR_SIZE` disk group attributes with the `CREATE DISKGROUP` and `ALTER DISKGROUP` SQL statements to specify a logical sector size and physical sector size for a disk group.

Note:

To specify the logical and physical sector size, the disks must support sector size emulation.

The `SECTOR_SIZE` disk group attribute can be set to 512, 4096, or 4K if the disks support those values. The default value is platform dependent. The `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` disk group attributes must be set to 11.2 or higher to enable the setting of a value for the sector size during disk group creation. The `COMPATIBLE.ASM` disk group attribute must be set to 12.2 or higher to alter the sector size.

The LOGICAL_SECTOR_SIZE disk group attribute can explicitly set the logical sector size of the disk group instead of the size that is estimated from the disks that join the disk group. The LOGICAL_SECTOR_SIZE disk group attribute can be set to 512, 4096, or 4K. The value specifies the smallest possible I/O that the disk group can accept, and also limits what types of disks can join the disk group. The COMPATIBLE.ASM disk group attribute must be set to 12.2 or higher to set the logical sector size during creation of a disk group or alter the size of a disk group that has been created.

Table 4-1 (page 4-16) lists the possible valid combinations of the logical sector size and sector size disk group attributes during disk group creation for disk sizes and types. The first column describes the user-specified values for logical sector size and sector size attributes during disk group creation. NS indicates that the no value was specified.

The top row of the table describes the combination of the disks used. A 512 native sector size disk has sector size equal to logical sector size equal to 512. The 512 native and 512 emulated heading describes the combination of 512 native disks and 512 emulated disk. The 512 emulated disk has a 4 k physical sector size and a 512 logical sector size. The values in each of the table cells represents the final value of sector sizes of the disk group. Error indicates that the disk group creation fails and an error message is displayed.

Table 4-1 Possible combinations of sector size and logical sector size by disk size and type

SECTOR_SIZE/ LOGICAL_SECT OR_SIZE specified at disk group creation	512 native disks	512 emulated disks	4K native disks	512 native disks + 512 emulated disks	512 emulated disks + 4K native disks
4K / 4K	Error	4K / 4K	4K / 4K	Error	4K / 4K
NS / 4K	Error	4K / 4K	4K / 4K	Error	4K / 4K
4K / 512	Error	4K / 512	Error	Error	Error
512 / 512	512 / 512	512 / 512	Error	512 / 512	Error
512 / NS	512 / 512	512 / 512	Error	512 / 512	Error
4K / NS	Error	4K / 512	4K / 4K	Error	4K / 4K
NS / 512	512 / 512	512 / 512	Error	512 / 512	Error
NS / NS	512 / 512	512 / 512	4K / 4K	512 / 512	4K / 4K

You can determine the logical sector and sector size values that have either been assumed or explicitly set for a successful disk group creation by querying V\$ASM views or running ASMCMD commands. The LOGICAL_SECTOR_SIZE and SECTOR_SIZE columns are included in the V\$ASM_DISKGROUP and V\$ASM_ATTRIBUTE views to display the values. These columns represent the logical sector size and sector size values of the disk group in bytes. The ASMCMD lsdg also displays the values of the logical sector size and sector size for the disk group. In addition, the disk group attributes are displayed with the V\$ASM_ATTRIBUTE view and the ASMCMD lsattr command. For example:

```
SQL> SELECT name, value FROM V$ASM_ATTRIBUTE
      WHERE (name = 'sector_size' OR name = 'logical_sector_size') AND group_number =
      1;
NAME                                VALUE
```



```

-----
sector_size                4096
logical_sector_size        512

SQL> SELECT logical_sector_size, sector_size FROM V$ASM_DISKGROUP
       WHERE group_number = 1;
LOGICAL_SECTOR_SIZE SECTOR_SIZE
-----
                    512        4096

```

As shown in [Example 4-3](#) (page 4-17), you can use the `LOGICAL_SECTOR_SIZE` and `SECTOR_SIZE` disk group attributes with the `CREATE DISKGROUP` and `ALTER DISKGROUP` SQL statements to specify the logical and physical sector size of the Oracle ASM disk group.

The following notes apply to the logical sector size and sector size disk group attributes.

- During the migration process from 512 bytes to 4096 bytes, password files stored on an Oracle ASM disk group do not need to be migrated or recreated. The password files are a special case and retain their 512 byte block size, even after the new sector sizes are 4096 bytes.
- If the value of the `LOGICAL_SECTOR_SIZE` column is 512, while the value in the `SECTOR_SIZE` column is 4096, then the disk group supports 512 sector size emulation. This means that I/O operations can be both 512 or 4096 in size. For example, we can have redo log files with a file block size of 512 bytes in the disk group that can be both read and written to, as well as creating new redo log files with a file block size of 4096 bytes.
- The physical sector size (`SECTOR_SIZE`) cannot be smaller than the logical sector size (`LOGICAL_SECTOR_SIZE`) in an Oracle ASM disk group.
- Not all disks support all of the possible `SECTOR_SIZE` values. The sector size attribute setting must be compatible with the physical hardware.

Example 4-3 Specifying logical and sector sizes for disks in a disk group

```

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'
  QUORUM FAILGROUP quorum_failgrp DISK '/devices/diskc1'
  ATTRIBUTE 'compatible.asm' = '12.2', 'compatible.rdbms' = '12.2',
            'sector_size'='4096', 'logical_sector_size'='512';

ALTER DISKGROUP data2
  SET ATTRIBUTE 'compatible.asm' = '12.2', 'compatible.rdbms' = '12.2';

ALTER DISKGROUP data2
  SET ATTRIBUTE 'sector_size'='4096', 'logical_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$ASM_ATTRIBUTE`, `V$ASM_DISK`, and `V$ASM_DISKGROUP` views
 - [ASMCMDB Disk Group Management Commands](#) (page 10-53) for information about ASMCMDB disk group management commands
-

4.2.5 Storing Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups

You can store Oracle Cluster Registry (OCR) and voting files in Oracle ASM disk groups.

You can also store a backup of the OCR file in a disk group. Voting files and OCR are two important components of Oracle Clusterware. Voting files manage information about node membership. OCR manages cluster and Oracle Real Application Clusters (Oracle RAC) database configuration information.

A quorum failure group is a special type of failure group that does not contain user data. Quorum failure groups are used for storing Oracle ASM metadata. A quorum failure group may also contain voting files if those files are stored in a disk group that contains a quorum failure group. Additionally, Oracle ASM uses a quorum failure group to help determine if the disk group can be mounted in the event of the loss of one or more failure groups.

Because disks in quorum failure groups do not contain user data, a quorum failure group is not considered when determining redundancy requirements in respect to storing user data.

You can manage and monitor OCR and voting files with `CRSCTL` and `ocrconfig` command-line tools, `CREATE` and `ALTER DISKGROUP` SQL statements, and `V$ASM` views.

The `CRSCTL` and `ocrconfig` commands enable the placement of OCR storage and Cluster Synchronization Services (CSS) voting files inside the disk groups managed by Oracle ASM.

The `QUORUM` and `REGULAR` keywords in `CREATE` and `ALTER DISKGROUP` SQL statements enable the specification of an additional qualifier for failure groups or disks when creating or altering a disk group. Note that `QUORUM` disks (or disks in `QUORUM` failure groups) cannot contain user data, while `REGULAR` disks (or disks in non-quorum failure groups) do not have this restriction.

You can use these keywords before the keyword `FAILGROUP` if a failure group is being explicitly specified. If the failure group is implicitly implied, you can use these keywords (`QUORUM`/`REGULAR`) 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 4-4](#) (page 4-19) 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 is required with a normal redundancy disk group. A QUORUM failure group is not considered when determining redundancy requirements in respect to storing user data. The COMPATIBLE .ASM disk group compatibility attribute must be set to 11.2 or greater to store OCR or voting files in a disk group.

The FAILGROUP_TYPE column in the V\$ASM_DISK or V\$ASM_DISK_STAT views specifies the 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 the V\$ASM_DISK or V\$ASM_DISK_STAT views specifies whether a disk contains a voting file. The value for this column is either Y if the disk contains a voting file or N if not. The value of USABLE_FILE_MB in V\$ASM_DISKGROUP or 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.

Example 4-4 Using the QUORUM keyword when creating a failure group

```
CREATE DISKGROUP mgmt_dg NORMAL REDUNDANCY
  FAILGROUP fg1 DISK '/devices/diskg1'
  FAILGROUP fg2 DISK '/devices/diskg2'
  QUORUM FAILGROUP fg3 DISK '/devices/diskg3'
  ATTRIBUTE 'compatible.asm' = '12.2.0.0.0';
```

See Also:

- [Oracle ASM Failure Groups](#) (page 4-40) for information about failure groups
 - [Views Containing Oracle ASM Disk Group Information](#) (page 6-1) for information about V\$ASM views
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for information about configuring voting files and Oracle Cluster Registry (OCR)
 - *Oracle Clusterware Administration and Deployment Guide* for information about managing Oracle Cluster Registry and voting files
 - *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
-

4.3 Altering Disk Groups

You can alter a disk group with SQL*Plus, ASMCA, or ASMCMD commands.

You can use the ALTER DISKGROUP SQL statement to alter a disk group configuration. You can also alter a disk group configuration with Oracle ASM Configuration Assistant (ASMCA) and the ASMCMD chdg command.

You can add, resize, or drop disks while the database remains online. Whenever possible, multiple operations in a single ALTER DISKGROUP statement are recommended. Grouping operations in a single ALTER DISKGROUP statement can reduce rebalancing operations.

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.

Use the `REBALANCE WAIT` clause to cause 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.

The following topics are discussed:

- [Managing Oracle ADVM Volumes in a Disk Group](#) (page 4-21)
- [Adding Disks to a Disk Group](#) (page 4-22)
- [Using the ALTER DISKGROUP SQL Statement to Add Disks to a Disk Group](#) (page 4-22)
- [Replacing Disks in Disk Groups](#) (page 4-24)
- [Renaming Disks in Disk Groups](#) (page 4-24)
- [Dropping Disks from Disk Groups](#) (page 4-25)
- [Intelligent Data Placement](#) (page 4-26)
- [Resizing Disks in Disk Groups](#) (page 4-27)
- [Undropping Disks in Disk Groups](#) (page 4-28)
- [Manually Rebalancing Disk Groups](#) (page 4-28)
- [Tuning Rebalance Operations](#) (page 4-31)
- [Scrubbing Disk Groups](#) (page 4-32)

See Also:

- [Managing Disk Groups with ASMCA](#) (page 9-2) for information about altering a disk group configuration with Oracle ASM Configuration Assistant (ASMCA)
 - [chdg](#) (page 10-54) for information about altering a disk group configuration with ASMCMD `chdg`
 - [Manually Rebalancing Disk Groups](#) (page 4-28) for more information about the `REBALANCE POWER` clause
 - *Oracle Database SQL Language Reference* for information about the `ALTER DISKGROUP` SQL statement
-

4.3.1 Managing Oracle ADVM Volumes in a Disk Group

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 be used to host an Oracle ACFS file system.

The compatibility attributes `COMPATIBLE . ASM` and `COMPATIBLE . ADVM` must be set to 11.2 or higher for the disk group.

The `ALTER DISKGROUP VOLUME` SQL statements enable you to manage Oracle ADVM volumes, including the functionality to add, modify, resize, disable, enable, and drop volumes. The following are examples of the `ALTER DISKGROUP VOLUME` statement.

If the volume is hosting an Oracle ACFS file system, then you cannot resize that volume with the SQL `ALTER DISKGROUP` statement. You must use the `acfsutil size` command instead.

Example 4-5 Managing volumes with `ALTER DISKGROUP VOLUME` statements

```
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.
```

See Also:

- [Disk Group Compatibility Attributes](#) (page 4-51) for information about disk group compatibility attributes
 - [acfsutil size](#) (page 16-164) for information about the resizing a volume with `acfsutil size`
 - [Overview of Oracle ASM Dynamic Volume Manager](#) (page 11-49) for information about Oracle ADVM
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about managing Oracle ADVM volumes with ASMCMD
 - [ASMCA GUI Tool for Managing Oracle ACFS and Oracle ADVM](#) (page 14-1) for information about managing Oracle ADVM volumes with ASMCA
 - *Oracle Database SQL Language Reference* for information about the ALTER DISKGROUP SQL statement
-

4.3.2 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.

You can use the same syntax with the ALTER DISKGROUP statement that you use to add a disk or failure group with the CREATE DISKGROUP statement. After you add new disks, the new disks gradually begin to accommodate their share of the workload as rebalancing progresses.

When adding a disk, the disk must be the same size as the other disks in the disk group.

See Also:

- [Example 4-1](#) (page 4-13) for an example of the CREATE DISKGROUP SQL statement
 - [Using the ALTER DISKGROUP SQL Statement to Add Disks to a Disk Group](#) (page 4-22) for an example of using SQL statements to add disks to a disk group
 - [Managing Disk Groups with ASMCA](#) (page 9-2) for information about adding disks to a disk group with ASMCA
 - *Oracle Database SQL Language Reference* for information about the ALTER DISKGROUP SQL statement
-

4.3.3 Using the ALTER DISKGROUP SQL Statement to Add Disks to a Disk Group

The SQL statements presented in the following 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 \(page 6-1\)"](#).

The following statement would fail because `/devices/diska1` through `/devices/diska4` currently belong to the disk group `data1`.

```

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

```

The following statement successfully adds 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 the disk as a member 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;
```

4.3.4 Replacing Disks in Disk Groups

A disk or multiple disks in a disk group can be replaced, rather than dropped and added back.

The single replace operation is more efficient than dropping and adding disks. This operation is especially useful when disks are missing or damaged.

For example, you can issue the following statement to replace the `diskc7` disk with another disk identified by the `/devices/diskc18` path.

```
SQL> ALTER DISKGROUP data2 REPLACE DISK diskc7 WITH '/devices/diskc18' POWER 3;
```

The `power` option operates the same as the `power` option for the `ALTER DISKGROUP REBALANCE` statement, except that the `power` option cannot be set to 0.

The `ALTER DISKGROUP SQL` statement with the `REPLACE` clause includes a `WAIT` or `NOWAIT` option, plus the `FORCE` option.

Note:

- [Manually Rebalancing Disk Groups](#) (page 4-28) for information about the `power` option in rebalancing a disk group
 - [ASM_POWER_LIMIT](#) (page 3-11) for information about the `ASM_POWER_LIMIT` initialization parameter
-
-

4.3.5 Renaming Disks in Disk Groups

You can rename a disk in a disk group with the `ALTER DISKGROUP RENAME DISK SQL` statement. In one SQL statement, you can rename one or multiple disks, or rename all disks in a disk group using the `RENAME DISKS ALL` clause.

For example, you can rename disks as follows:

```
SQL> ALTER DISKGROUP fra2 MOUNT RESTRICTED;
```

```
SQL> ALTER DISKGROUP fra2 RENAME DISK 'FRA1_0001' TO 'FRA2_0001',
    'FRA1_0002' TO 'FRA2_0002';
```

The `ALTER DISKGROUP diskgroupname RENAME DISKS ALL` statement can be run after the `renamedg` utility to change the names of the disks in the renamed disk group. For information about `renamedg`, refer to "[Renaming Disk Groups](#) (page 4-70)".

When you run the `ALTER DISKGROUP diskgroupname RENAME DISKS ALL` statement, any disk name that is not in the format `diskgroupname_number` is renamed to that format. Disk names that are already in the `diskgroupname_number` format are not changed.

You can only use the `RENAME DISK` operation when the disk group that contains the disk is in the `MOUNT RESTRICTED` state. If any disks in the disk group are offline, then the `RENAME` operation fails. If the new disk name exists, then the `RENAME` operation fails. You must have `SYSASM` privileges to rename a disk.

Note:

If a disk is labeled by any subsystem; such as `ASMLIB`, `ASMFD`, `EXADATA`, or `ASMTOOL`; then that label becomes the disk name by design. The `ALTER DISKGROUP RENAME DISKS SQL` statement on these subsystems has no effect and the disk is not renamed.

4.3.6 Dropping Disks from Disk Groups

You can drop disks from a disk group with 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 one time and ensuring that there is enough space for the rebalance operation to succeed.

Caution:

The `ALTER DISKGROUP . . . DROP DISK SQL` statement returns to SQL prompt 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.

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. Use `DROP FORCE` with caution.

The statements in [Example 4-6](#) (page 4-26) demonstrate how to drop disks from the disk group `data1`. The first example drops `diska5` from disk group `data1`. The second example drops `diska5` from disk group `data1`, and also illustrates how multiple actions are possible with one `ALTER DISKGROUP` statement.

Example 4-6 Dropping disks from disk groups

```
ALTER DISKGROUP data1 DROP DISK diska5;

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

See Also:

- [Managing Disk Groups with ASMCA](#) (page 9-2) for information about dropping disks from a disk group with ASMCA
 - [Using the ALTER DISKGROUP SQL Statement to Add Disks to a Disk Group](#) (page 4-22) for an example of the `ALTER DISKGROUP` statement
 - *Oracle Database Reference* for information about the `V$ASM_OPERATION` view
 - *Oracle Database SQL Language Reference* for information about the `ALTER DISKGROUP` SQL statement
-
-

4.3.7 Intelligent Data Placement

Note:

The Intelligent Data Placement (IDP) feature has been deprecated in Oracle ASM 12c Release 2 (12.2) and may be desupported in a future release.

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, reducing latency. Intelligent Data Placement also enables 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
  ATTRIBUTE (
    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'
  ATTRIBUTE (
    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.

Oracle ASM Configuration Assistant (ASMCA) supports Intelligent Data Placement with template creation during disk group alterations.

See Also:

- [Manually Rebalancing Disk Groups](#) (page 4-28) for information about rebalance operations
 - [Managing Disk Groups with ASMCA](#) (page 9-2) for information about disk group alterations with ASMCA
 - [Viewing Disk Region Information](#) (page 6-6) and [iostat](#) (page 10-58) for information about displaying information about Intelligent Data Placement regions
-
-

4.3.8 Resizing Disks in Disk Groups

The `RESIZE` clause of `ALTER DISKGROUP` enables you to resize disks in a disk group.

When resizing disks in a disk group, all the disks must be of equal size.

If you do not specify a new size in the `SIZE` clause, then Oracle ASM uses the size of the disks as returned by the operating system. The new size is written to the Oracle ASM disk header and if the size 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.

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

Example 4-7 Resizing disks in disk groups

```
ALTER DISKGROUP data1 RESIZE ALL SIZE 100G;
```

4.3.9 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.

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

Example 4-8 Undropping disks in disk groups

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

Note:

- *Oracle Database SQL Language Reference* for information about the `DROP DISKGROUP` SQL statement
 - *Oracle Database SQL Language Reference* for information about the `ALTER DISKGROUP` SQL statement
-
-

4.3.10 Manually Rebalancing Disk Groups

You can manually rebalance the files in a disk group using the `REBALANCE` clause of the `ALTER DISKGROUP` statement.

Manual rebalancing would normally not be required because Oracle ASM automatically rebalances disk groups when their configuration changes, including changes to file groups. However, you might want to do a manual rebalance operation 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 the speed of the rebalance operation. A higher value increases the speed of the rebalance operation. It can be set to a minimum value of 0 which stops a rebalancing operation until the statement is either implicitly or explicitly rerun.

The default rebalance power is set by the `ASM_POWER_LIMIT` initialization parameter. The range of values for the `POWER` clause is the same for the `ASM_POWER_LIMIT` initialization parameter.

The power level of an ongoing rebalance operation can be changed by entering the rebalance statement using the `MODIFY POWER` clause. For example, the following SQL statement modifies the power setting to 10 of an ongoing rebalance.

```
ALTER DISKGROUP data2 REBALANCE MODIFY POWER 10;
```

The the following SQL statement modifies the power setting to the default value.

```
ALTER DISKGROUP data2 REBALANCE MODIFY POWER;
```

When you rebalance a disk group, you have the option choosing the following phase options using the `WITH` or `WITHOUT` keywords, except where noted:

- `RESTORE`

The `RESTORE` phase includes `RESYNC`, `RESILVER`, and `REBUILD` operations. The `RESTORE` phase is always run by default and cannot be excluded.

 - `RESYNC`

This operation synchronizes the stale extents on the disks that are being brought online.
 - `RESILVER`

This operation is specific only to Exadata systems. During this phase, data is copied from one mirror to the mirror with stale data.
 - `REBUILD`

This operation restores the redundancy of forcing disks only. Forcing disks are those disks that have been dropped with the `force` option.
- `BALANCE`

This phase restores redundancy of all the disks in the disk group, including file groups, and also balances extents on Oracle ASM disks.
- `PREPARE`

This phase completes the work corresponding to the prepare SQL operation. This phase is enabled only for `FLEX` or `EXTENDED` redundancy disk groups. `COMPATIBLE .ASM` must be set to 12.2 or higher.
- `COMPACT`

This phase defragments and compacts extents across Oracle ASM disks.

The rebalance operation executes the phases in the optimal order. If none of the phases are specified with the `ALTER DISKGROUP . . . REBALANCE` SQL statement, then the behavior remains the same as the rebalance operation previous to the Oracle ASM 12.2 release during which all rebalance phases are run.

The following examples show the use of the phase options with `ALTER DISKGROUP . . . REBALANCE`:

```
ALTER DISKGROUP data2 REBALANCE WITH BALANCE COMPACT;
```

```
ALTER DISKGROUP data3 REBALANCE WITHOUT BALANCE;
```

The `ALTER DISKGROUP . . . REBALANCE` command returns immediately by default so that you can run other commands while the rebalance operation takes place asynchronously in the background. You can query the `V$ASM_OPERATION` view for the status of the rebalance operation. When not all rebalance phases are run, the `V$ASM_OPERATION` view displays an `ORA-59000` error at the end of the rebalance.

To cause the `ALTER DISKGROUP . . . REBALANCE` command to wait until the rebalance operation is complete before returning, add the `WAIT` keyword to the `REBALANCE` clause. The wait functionality 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 continues 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. If you have initiated multiple rebalances on different disk groups on a single node, then Oracle processes these operations in parallel on additional nodes if available; otherwise the rebalances are performed serially on the single node. You can explicitly 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 a manual rebalance may be required.

The following example manually rebalances the disk group `data2`. The command does not return until the rebalance operation is complete. All supported phases of the rebalance operation are executed. Note that specifying `RESTORE` is not necessary as that phase is always run during a rebalance.

Example 4-9 Manually rebalancing a disk group

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

See Also:

- [ASM_POWER_LIMIT](#) (page 3-11) for information about the `ASM_POWER_LIMIT` initialization parameter
 - [Tuning Rebalance Operations](#) (page 4-31) for information about rebalancing operations
 - *Oracle Database SQL Language Reference* for information about the `ALTER DISKGROUP SQL` statement
-
-

4.3.11 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 range of values for the `POWER` clause is the same for the `ASM_POWER_LIMIT` initialization parameter.

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.

Oracle ASM tries to keep a rebalance I/O for each unit of power. Each I/O requires PGA memory for the extent involved in the relocation.

The default value of 1 minimizes disruption to other applications. The appropriate value is dependent on your hardware configuration, performance requirements, and availability requirements. Oracle ASM always adjusts the power to fit available memory.

When the `COMPATIBLE . ASM` disk group is set to 11.2.0.2 or higher, the rebalance operation may be run as one process using asynchronous I/O. You can check the Oracle ASM alert log for details on the rebalance process.

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.

You can also affect rebalance behavior with the `CONTENT . CHECK` and `THIN_PROVISIONED` disk group attributes.

The `EXPLAIN WORK SQL` statement determines the amount of work for a rebalance operation and the resulting calculations are displayed in the `V$ASM_ESTIMATE` view.

For example:

```
EXPLAIN WORK FOR ALTER DISKGROUP data DROP DISK data_0000;
Explained.
```

```
SELECT est_work FROM V$ASM_ESTIMATE;
EST_WORK
-----
      4211
```

```
EXPLAIN WORK SET STATEMENT_ID='online' FOR ALTER DISKGROUP data
  ONLINE disk data_000;
Explained.
```

```
SELECT est_work FROM V$ASM_ESTIMATE WHERE STATEMENT_ID='online';
EST_WORK
-----
      421
```

The `V$ASM_ESTIMATE` view provides information for adjusting `ASM_POWER_LIMIT` and the resulting power of rebalance operations. The `EST_WORK` column provides an estimate of the number of allocation units that have to be moved by the rebalance operation to complete.

The `PASS` column of `V$ASM_OPERATION` is updated for resync and rebalance operations. The contents of the column can be `RESYNC`, `REBALANCE`, or `COMPACT`. For

example, the following SQL query shows values in the `PASS` column during a rebalance operation.

```
SQL> SELECT GROUP_NUMBER, PASS, STATE FROM V$ASM_OPERATION;
```

GROUP_NUMBER	PASS	STAT
2	RESYNC	WAIT
2	REBALANCE	WAIT
2	COMPACT	WAIT

See Also:

- [ASM_POWER_LIMIT](#) (page 3-11) for information about the `ASM_POWER_LIMIT` initialization parameter
 - [Summary of Disk Group Attributes](#) (page 4-3) for information about disk group attributes
 - [Manually Rebalancing Disk Groups](#) (page 4-28) for more information about rebalancing disk groups
-

4.3.12 Scrubbing Disk Groups

Oracle ASM disk scrubbing improves availability and reliability by searching for data that may be less likely to be read. Disk scrubbing checks logical data corruptions and repairs them automatically in normal and high redundancy disks groups. The scrubbing process repairs logical corruptions using the mirror disks. Disk scrubbing can be combined with disk group rebalancing to reduce I/O resources. The disk scrubbing process has minimal impact to the regular I/O in production systems.

You can perform scrubbing on a disk group, a specified disk, or a specified file of a disk group with the `ALTER DISKGROUP` SQL statement. For example, the following SQL statements show various options used when running the `ALTER DISKGROUP disk_group SCRUB` SQL statement.

```
SQL> ALTER DISKGROUP data SCRUB POWER LOW;
```

```
SQL> ALTER DISKGROUP data SCRUB FILE '+DATA/ORCL/DATAFILE/example.266.806582193'  
REPAIR POWER HIGH FORCE;
```

```
SQL> ALTER DISKGROUP data SCRUB DISK DATA_0005 REPAIR POWER HIGH FORCE;
```

When using `ALTER DISKGROUP` with the `SCRUB` option, the following items apply:

- The optional `REPAIR` option automatically repairs disk corruptions. If the `REPAIR` option is not specified, then the `SCRUB` option only checks and reports logical corruptions of the specified target.
- The optional `POWER` value can be set to `AUTO`, `LOW`, `HIGH`, or `MAX`. If the `POWER` option is not specified, the power value defaults to `AUTO` and the power adjusts to the optimum level for the system.
- If the optional `WAIT` option is specified, the command returns after the scrubbing operation has completed. If the `WAIT` option is not specified, the scrubbing operation is added into the scrubbing queue and the command returns immediately.

- If the optional `FORCE` option is specified, the command is processed even if the system I/O load is high or scrubbing has been disabled internally at the system level.

The scrubbing process is visible in the fields of the `V$ASM_OPERATION` view while the scrubbing process is running.

4.4 Oracle ASM Disk Discovery

Disk discovery is the mechanism used to find the operating system names for disks Oracle ASM can access.

This mechanism finds 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](#) (page 4-33)
- [Disk Discovery Rules](#) (page 4-34)
- [Improving Disk Discovery Time](#) (page 4-35)

See Also:

[ASM_DISKSTRING](#) (page 3-10) for additional information about disk discovery and the `ASM_DISKSTRING` initialization parameter

4.4.1 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 ASM Configuration Assistant (ASMCA) operations or `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. For example, on Linux computers, the administrator might have used Oracle ASM Filter Driver or ASMLIB to prepare the disk for Oracle ASM.

[Example 4-10](#) (page 4-34) shows a SQL query on `V$ASM_DISK` that displays the header status of a group of disks.

Example 4-10 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

4.4.2 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, if 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](#) (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.

4.4.3 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 may 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 Oracle ASM Filter Driver or a third-party vendor, then you may 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 Also:

- [Administering Oracle ASM Filter Driver](#) (page 20-1) for information about Oracle ASM Filter Driver
 - [Oracle ASM and Multipathing](#) (page 2-3) and your multipathing vendor documentation for information about multipathing
-
-

4.5 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. For unprotected files, data extents could be completely missing. Other causes of

reduced redundancy files are disks running out of space or an insufficient number of failure groups.

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:

Note:

The values in the `REQUIRED_MIRROR_FREE_MB` and `USABLE_FILE_MB` columns of the `V$ASM_DISKGROUP` view are 0 if the disk group type is `FLEX` or `EXTEND`. The 0 value in this case means no value is given, not a zero amount in the `REQUIRED_MIRROR_FREE_MB` and `USABLE_FILE_MB` columns.

- `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 requirement ensures 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:

- A normal redundancy disk group with more than two `REGULAR` 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.

- A high redundancy disk group with more than three `REGULAR` failure groups

The value is the total raw space for all of the disks in the two largest failure groups.

See Also:

Oracle Exadata documentation for information about calculating `REQUIRED_MIRROR_FREE_MB` on Oracle Exadata systems.

Oracle Database Appliance documentation for information about calculating `REQUIRED_MIRROR_FREE_MB` on Oracle Database Appliance systems.

- `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. For information about negative values of `USABLE_FILE_MB`, refer to "[Negative Values of USABLE_FILE_MB](#) (page 4-38)".
- `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. There may be situations where the value in the `FREE_MB` column shows unused capacity but because one Oracle ASM disk is full, database writes fail because of the imbalance in the disk group. Ensure 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. 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:

$$(FREE_MB - REQUIRED_MIRROR_FREE_MB) / 2 = USABLE_FILE_MB$$

$$(3768 - 1024) / 2 = 2744 / 2 = 1372$$

4.5.1 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.

4.6 Oracle ASM Mirroring and Disk Group Redundancy

This section contains the following topics:

- [Mirroring, Redundancy, and Failure Group Options](#) (page 4-38)
- [Oracle ASM Recovery from Read and Write I/O Errors](#) (page 4-42)
- [Oracle ASM Fast Mirror Resync](#) (page 4-43)
- [Preferred Read Failure Groups](#) (page 4-46)

4.6.1 Mirroring, Redundancy, and Failure Group Options

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 to normal, high, flex, and extended redundancy disk groups. You can define the failure groups for each disk group when you create or alter the disk group.

There are multiple types of disk groups based on the Oracle ASM redundancy level. [Table 4-2](#) (page 4-38) 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 4-2 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
Flex redundancy	Two-way, three-way, unprotected (none)	Two-way (newly-created)

Table 4-2 (Cont.) Mirroring options for Oracle ASM disk group types

Disk Group Type	Supported Mirroring Levels	Default Mirroring Level
Extended redundancy	Two-way, three-way, unprotected (none)	Two-way

For normal and high disk group types, 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.

For the flex group type, the number of failures tolerated before dismount demands on the number of failure groups. For five or more failure groups, two disk failures are tolerated. For three or four failure groups, one disk failure is tolerated.

For the extended disk group type, each site is similar to a flex disk group. If the site has five failure groups or more, two disk failures within a site can be tolerated before the site becomes compromised. If the site has three or four failure groups, the site can tolerate one disk failure before the site is compromised. When two sites are compromised, the disk group dismounts. An extended disk group requires a minimum of three failure groups for each data site.

For flex and extended disk groups, mirroring describes the availability of the files within a disk group, not the disk group itself. For example: If a file is unprotected in a flex disk group that has five failure groups, then after one failure the disk group is still mounted, but the file becomes unavailable.

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 causes 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.

A file specified with HIGH redundancy (three-way mirroring) in a NORMAL redundancy disk group provides additional protection from a bad disk sector in one disk, plus the failure of another disk. However, this scenario does not protect against the failure of two disks.
- High redundancy

Oracle ASM provides three-way (triple) mirroring by default. A loss of two Oracle ASM disks in different failure groups is tolerated.
- Flex redundancy

Oracle ASM provides two-way mirroring by default for newly-created flex disk groups. For migrated flex disk groups, the default values are obtained from the template values in the normal or high redundancy disk groups before migration. For migration from normal redundancy, if the template defaults were not changed, then the flex defaults are two-way mirroring. For migration from high redundancy, if the template defaults were not changed, then the flex defaults are three-way mirroring.

- Extended redundancy
Oracle ASM provides two-way mirroring by default. The redundancy setting describes redundancy within a data site. For example: If there is a two-way mirrored file in a two-data-site extended disk group, then there are four copies of the file, two in each data site.

Oracle ASM file groups in a flex or extended disk group can have different redundancy levels.

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.

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](#) (page 4-40)
- [How Oracle ASM Manages Disk Failures](#) (page 4-41)
- [Guidelines for Using Failure Groups](#) (page 4-41)
- [Failure Group Frequently Asked Questions](#) (page 4-42)

See Also:

- [Managing Oracle ASM Flex Disk Groups](#) (page 4-72) for information about Oracle ASM flex disk groups, extended disk groups, and file groups
 - [Managing Disk Group Templates](#) (page 5-19) for information about specifying Oracle ASM disk group templates
-
-

4.6.1.1 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 more 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.

Oracle recommends a minimum of three failure groups for normal redundancy disk groups and five failure groups for high redundancy disk groups to maintain the necessary number of copies of the Partner Status Table (PST) and to ensure robustness with respect to storage hardware failures.

In the event of a system failure, three failure groups in a normal redundancy disk group allow a comparison among three PSTs to accurately determine the most up to date and correct version of the PST, which could not be done with a comparison between only two PSTs. Similarly with a high redundancy disk group, if two failure groups are offline, then Oracle ASM would be able to make a comparison among the three remaining PSTs.

If configuring an extra failure group presents a problem with storage capacity management, then a quorum failure group can be used as the extra failure group to store a copy of the PST. A quorum failure group does not require the same capacity as the other failure groups.

Failure groups can be specified as regular or quorum failure groups. For information about quorum failure groups, see "[Storing Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#) (page 4-18)".

See Also:

Oracle Exadata documentation for information about Oracle Exadata failure groups

4.6.1.2 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.

4.6.1.3 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.

4.6.1.4 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?](#) (page 4-42)
- [How are Multiple Failure Groups Recovered after Simultaneous Failures?](#) (page 4-42)
- [When Should External, Normal, or High Redundancy Be Used?](#) (page 4-42)

How Many Failure Groups Should I Create?

Choosing the number of failure groups to create depends on the types of failures that must 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 deployment of additional layers of software complexity in your Oracle Database environment.

4.6.2 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 can 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, the 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 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](#) (page 10-75)".

4.6.3 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 disk 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.

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 and is only applicable to normal and high redundancy disk groups.

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, ensure 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 during which you are able 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_TIMER` 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 the `ALTER DISKGROUP OFFLINE DISK` 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 specified for the `DISK_REPAIR_TIME` attribute for the disk group is used.

Altering the `DISK_REPAIR_TIME` attribute has no effect on offline disks. The new value is used for any disks that go offline after the attribute is updated. 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 ONLINE DISK`. 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. 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 disks 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 run ALTER DISKGROUP ONLINE statements displays the name and state of the current operation that you are performing. For example, the following SQL query shows values in the PASS column during an online operation.

```
SQL> SELECT GROUP_NUMBER, PASS, STATE FROM V$ASM_OPERATION;
```

GROUP_NUMBER	PASS	STAT
1	RESYNC	RUN
1	REBALANCE	WAIT
1	COMPACT	WAIT

An offline operation does not generate a display in a V\$ASM_OPERATION view query.

You can set the FAILGROUP_REPAIR_TIME and CONTENT . TYPE disk group attributes. The FAILGROUP_REPAIR_TIME disk group attribute specifies a default repair time for the failure groups in the disk group. The CONTENT . TYPE disk group attribute specifies the type of data expected to be stored in a disk group. You can set these attributes with ASMCA, ASMCMD mkd_g, or SQL CREATE and ALTER DISKGROUP statements.

The ASMCMD l_{sop} command shows the resync time estimate. There are separate rows in the V\$ASM_OPERATION table for different phases of rebalance: disk resync, rebalance, and data compaction.

The ASMCMD on_{line} command has a power option to specify the power for the online operation. The SQL ALTER DISKGROUP REPLACE DISK statement also has the power option.

The ASMCMD `chdg` command provides the `replace` option in addition to the `add` and `drop` tags. The ASMCMD `mkdgm` command has an additional time parameter (`-t`) to specify the time to offline a failure group.

See Also:

- [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - [Managing Disk Group Attributes](#) (page 4-2) for information about disk group attributes
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* and [Upgrading and Patching Oracle ASM](#) (page 3-25) for information about upgrading Oracle ASM
 - *Oracle Database SQL Language Reference* for information about ALTER DISKGROUP
 - *Oracle Database SQL Language Reference* for information about CREATE DISKGROUP
-

4.6.4 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 the preferred read failure groups feature 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.

Note:

In an Oracle extended cluster, which contains nodes that span multiple physically separated sites, the `PREFERRED_READ.ENABLED` disk group attribute controls whether preferred read functionality is enabled for a disk group. If preferred read functionality is enabled, then this functionality enables an instance to determine and read from disks at the same site as itself, which can improve performance. Whether or not `PREFERRED_READ.ENABLED` has been enabled, preferred read can be set at the failure group level on an Oracle ASM instance or a client instance in a cluster with the `ASM_PREFERRED_READ_FAILURE_GROUPS` initialization parameter, which is available for backward compatibility.

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.

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 are multiple mirrored copies 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 multiple failure groups on one site can cause the loss of access to the disk group by the other sites if the site containing multiple failure groups fails. In addition, by having multiple failure groups 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.

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:

If you do not specify failure groups for a disk group, each disk in the disk group belongs to its own failure group. Oracle does not recommend that you configure multiple preferred read failure groups in a disk group for an Oracle ASM instance. For any given instance, if you specify multiple failure groups in the same disk group as preferred read, a warning message is written to the alert log.

See Also:

- [PREFERRED_READ.ENABLED](#) (page 4-7) for information about the `PREFERRED_READ.ENABLED` disk group attribute
 - [ASM_PREFERRED_READ_FAILURE_GROUPS](#) (page 3-12) for more information about the `ASM_PREFERRED_READ_FAILURE_GROUPS` initialization parameter
 - *Oracle Real Application Clusters Administration and Deployment Guide* for information about configuring preferred read disks in extended clusters
 - *Oracle Database Reference* for details about the `V$ASM*` dynamic performance views
-
-

4.7 Performance and Scalability Considerations for Disk Groups

This section discusses the following considerations for evaluating disk group performance:

- [Determining the Number of Disk Groups](#) (page 4-48)
- [Performance Characteristics When Grouping Disks](#) (page 4-48)
- [Oracle ASM Storage Limits](#) (page 4-49)

4.7.1 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 should a disk group failure occur.

For information about creating disk groups for a new installation, refer to "[Creating Disk Groups for a New Oracle Installation](#) (page 4-14)".

4.7.2 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 characteristics coexist 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.

4.7.3 Oracle ASM Storage Limits

Oracle ASM provides near unlimited capacity for future growth, but does have some storage limits.

Oracle ASM has the following limits on the number of disk groups, disks, and files:

- 511 disk groups in a storage system for Oracle Database 12c Release 1 or later
- 10,000 Oracle ASM disks in a storage system
- 1 million files for each disk group

Without any Oracle Exadata Storage, Oracle ASM has the following storage limits if the `COMPATIBLE.ASM` or `COMPATIBLE.RDBMS` disk group attribute is set to less than 12.1:

- 2 terabytes (TB) maximum storage for each Oracle ASM disk
- 20 petabytes (PB) maximum for the storage system

Without any Oracle Exadata Storage, Oracle ASM has the following storage limits if the `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` disk group attributes are set to 12.1 or greater:

- 4 PB maximum storage for each Oracle ASM disk with the allocation unit (AU) size equal to 1 MB
- 8 PB maximum storage for each Oracle ASM disk with the AU size equal to 2 MB
- 16 PB maximum storage for each Oracle ASM disk with the AU size equal to 4 MB
- 32 PB maximum storage for each Oracle ASM disk with the AU size equal to 8 MB
- 320 exabytes (EB) maximum for the storage system

With all Oracle Exadata Storage, Oracle ASM has the following storage limits:

- 4 PB maximum storage for each Oracle ASM disk with the AU size equal to 1 MB
- 8 PB maximum storage for each Oracle ASM disk with the AU size equal to 2 MB
- 16 PB maximum storage for each Oracle ASM disk with the AU size equal to 4 MB
- 32 PB maximum storage for each Oracle ASM disk with the AU size equal to 8 MB
- 320 EB maximum for the storage system

The maximum size limit of a disk group equals the maximum disk size multiplied by the maximum number of disks in a disk group (10,000).

The maximum number of disks across all disk groups is 10,000. The 10,000 disks can be in one disk group or distributed across a maximum of 511 disk groups. This is a limitation on the number of Oracle ASM disks, not necessarily the number of spindles. A storage array could group multiple spindles into a LUN that is used as a single Oracle ASM disk.

File size limits are dependent on the value of the disk group compatibility attributes. Oracle ASM supports file sizes greater than 128 TB in any redundancy mode when the `COMPATIBLE.RDBMS` disk group attribute is set greater than 10.1.

If `COMPATIBLE.RDBMS` is set to 10.1, the file size limits are less. For example, with `COMPATIBLE.RDBMS` equal to 10.1 and the AU size equal to 1 MB, Oracle ASM file size limits are:

- External redundancy: 16 TB
- Normal redundancy: 5.8 TB
- High redundancy: 3.9 TB

Note:

Oracle Database supports data file sizes up to 128 TB depending on the file system. In addition, Oracle Database has a file size limit that is dependent on the `DB_BLOCK_SIZE` initialization parameter.

See Also:

- [Disk Group Compatibility](#) (page 4-50) for information about Oracle ASM disk group compatibility attributes
 - [Considerations When Setting Disk Group Compatibility](#) (page 4-60) for additional information about Oracle ASM file size limits
 - [Oracle ACFS Disk Space Usage](#) (page 17-2) for information about Oracle ACFS space usage
-
-

4.8 Disk Group Compatibility

This section describes disk group compatibility under the following topics:

- [Overview of Disk Group Compatibility](#) (page 4-50)
- [Disk Group Compatibility Attributes](#) (page 4-51)
- [Setting Disk Group Compatibility Attributes](#) (page 4-53)
- [Features Enabled By Disk Group Compatibility Attribute Settings](#) (page 4-57)
- [Considerations When Setting Disk Group Compatibility](#) (page 4-60)

4.8.1 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 12.2 can take advantage of Oracle ASM 12c

Release 2 (12.2) features. For information about the features enabled for combinations of compatibility attribute settings, refer to [Table 4-4](#) (page 4-58).

The disk group compatibility feature also enables environments to interoperate when they use disk groups from both Oracle Database 10g and Oracle Database 11g instances. For example, disk group compatibility settings that are set to Oracle Database 10g Release 2 (10.2) enable an Oracle Database 10g Release 2 (10.2) client to access a disk group created with Oracle ASM 12c Release 1 (12.1).

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 12.1, and RDBMS compatibility is set to 11.2, then the Oracle ASM software version must be at least 12.1, and the Oracle Database client software version must be at least 11.2. The `COMPATIBLE.ADVM` attribute determines whether the Oracle ASM Dynamic Volume Manager feature can create a volume in a disk group.

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. For examples of valid combinations of compatible attribute settings, refer to [Table 4-3](#) (page 4-53).

You can also set and update disk group attributes with the `ASMCMD setattr` command. For information about the `ASMCMD setattr` command, refer to [setattr](#) (page 10-75).

Note:

- To use a disk group with Oracle ASM 12c Release 2 (12.2), `COMPATIBLE.ASM` must be set to 11.2.0.2.0 or higher and `COMPATIBLE.RDBMS` must be set to 11.2 or higher.
 - The compatibility attributes are not set to zero releases, such as 11.0 or 12.0 or any value with all zeroes following the first dot.
 - 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. For more information, refer to [Reverting Disk Group Compatibility](#) (page 4-55).
 - 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.
-
-

4.8.2 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](#) (page 4-52)

- [COMPATIBLE.RDBMS](#) (page 4-52)
- [COMPATIBLE.ADVM](#) (page 4-52)

4.8.2.1 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 12c Release 2 (12.2), 12.2 is the default setting for the `COMPATIBLE.ASM` attribute when using ASMCA. The default setting is 10.1 when using the SQL `CREATE DISKGROUP` statement and the `ASMCMD mkdg` command.

4.8.2.2 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 12.1 or 12.2, then `COMPATIBLE.RDBMS` can be set to any value between 10.1 and 12.1 inclusively.

For Oracle ASM 12c, 10.1 is the default setting for the `COMPATIBLE.RDBMS` attribute when using the SQL `CREATE DISKGROUP` statement, the `ASMCMD mkdg` command, and ASMCA.

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.

4.8.2.3 COMPATIBLE.ADVM

The value for the disk group `COMPATIBLE.ADVM` attribute determines whether the disk group can contain Oracle ADVM 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 in the supported environment.

The default value of the `COMPATIBLE.ADVM` attribute depends on the level of the Oracle ASM release.

For more information about Oracle ADVM, see "[Overview of Oracle ASM Dynamic Volume Manager](#) (page 11-49)".

4.8.3 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](#) (page 4-53)
- [Using CREATE DISKGROUP with Compatibility Attributes](#) (page 4-54)
- [Using ALTER DISKGROUP with Compatibility Attributes](#) (page 4-54)
- [Viewing Compatibility Attribute Settings](#) (page 4-55)
- [Reverting Disk Group Compatibility](#) (page 4-55)

You can also set the disk group compatibility settings with Oracle ASM command-line utility (ASMCMD) and Oracle ASM Configuration Assistant (ASMCA). Refer to [Managing Oracle ASM With ASMCA](#) (page 9-1) and [Managing Oracle ASM with ASMCMD](#) (page 10-1).

Note:

Advancing the values for disk group compatibility attributes is an irreversible operation. See "[Reverting Disk Group Compatibility](#) (page 4-55)".

See Also:

Oracle Database SQL Language Reference for more information about the disk group compatibility SQL statements

4.8.3.1 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 '11.2' or '12.1'; Oracle assumes that any missing version number digits are zeros.

[Table 4-3](#) (page 4-53) shows some valid combinations of the disk group compatibility attributes and the valid Oracle ASM and database instance versions for each combination.

Table 4-3 Examples of disk group compatibility attribute settings

COMPATIBLE.ASM	COMPATIBLE.RDBMS	COMPATIBLE.ADV	Oracle 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	11.2	11.2	>= 11.2	>= 11.2

Table 4-3 (Cont.) Examples of disk group compatibility attribute settings

COMPATIBLE.ASM	COMPATIBLE.RDBMS	COMPATIBLE.ADVM	Oracle ASM Instance Version	COMPATIBLE Setting for RDBMS Instance
12.1	12.1	12.1	>= 12.1	>= 12.1
12.2	12.1	12.2	>= 12.2	>= 12.1

These are some possible combinations of Oracle ASM and database releases:

- The database release is Oracle Database 11g Release 2 (11.2) (database initialization parameter `COMPATIBLE` is set to 11.2) and the Oracle ASM release is Oracle ASM 11g Release 2 (11.2). The `COMPATIBLE.ASM` and `COMPATIBLE.RDBMS` disk group attributes are set to 10.2 for a disk group.

The Oracle ASM disk group functionality remains at Oracle Database 10g Release 2 (10.2).

- The database release is Oracle Database 10g Release 1 (10.1) and the Oracle ASM release is 11.2. `COMPATIBLE.ASM` is set to 11.2 and `COMPATIBLE.RDBMS` is set to 10.1 for a disk group.

The Oracle ASM features enabled for this combination of attribute settings are shown in [Features Enabled By Disk Group Compatibility Attribute Settings](#) (page 4-57).

- The database release is Oracle Database 12c Release 2 (12.2) (database initialization parameter `COMPATIBLE` is set to 12.2) and the Oracle ASM release is Oracle ASM 12c Release 2 (12.2). All the disk group compatibility attributes are set to 12.2 for a disk group.

The Oracle ASM features enabled for this combination of attribute settings are shown in [Features Enabled By Disk Group Compatibility Attribute Settings](#) (page 4-57).

4.8.3.2 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 12.1 and the RDBMS compatibility set to the default (the `COMPATIBLE.RDBMS` default is less than or equal to 12.1):

```
CREATE DISKGROUP data1 DISK '/dev/sd*'
  ATTRIBUTE 'compatible.asm' = '12.1';
```

The following example creates a normal redundancy disk group `data2` with the ASM, RDBMS, and ADVM compatibility set to 12.1:

```
CREATE DISKGROUP data2 DISK '/dev/sd*'
  ATTRIBUTE 'compatible.asm' = '12.1', 'compatible.rdbms' = '12.1',
  'compatible.advm' = '12.1';
```

4.8.3.3 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 4-3](#) (page 4-53). 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 12.1. An Oracle ASM instance must be at Oracle ASM 12c Release 1 (12.1) or higher to access the data3 disk group.

```
ALTER DISKGROUP data3 SET ATTRIBUTE 'compatible.asm' = '12.1';
```

The following example advances the COMPATIBLE.RDBMS and COMPATIBLE.ADVM compatibility attributes of the disk group data3 to 12.1. This example assumes that the value of COMPATIBLE.ASM is set to 12.1.

```
ALTER DISKGROUP data3 SET ATTRIBUTE 'compatible.rdbms' = '12.1',
ALTER DISKGROUP data3 SET ATTRIBUTE 'compatible.advm' = '12.1';
```

4.8.3.4 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.

You can also display the disk group compatibility attributes with the ASMCMD lsattr command.

See Also:

- [Example 6-1](#) (page 6-3) for an example querying the V\$ASM_ATTRIBUTE view
 - [lsattr](#) (page 10-60) for information about the ASMCMD lsattr command
 - *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
-

4.8.3.5 Reverting Disk Group Compatibility

You cannot revert disk group compatibility to a previous values. Instead you must create a new disk group.

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/asmspfile.ora' FROM SPFILE;
```
2. If the OCR and voting files are in the disk group, move them out of this disk group.
3. Back up any files that must be saved.
 - 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.
5. Move or restore the database files into the newly created disk group using either the `ALTER DATABASE MOVE DATAFILE` SQL statement or Recovery Manager (RMAN).
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.

See Also:

- [Using the CREATE DISKGROUP SQL Statement](#) (page 4-10) for information about creating a disk group
- [Moving Data Files Between Disk Groups With ALTER DATABASE](#) (page 5-11) for information about moving data files between disk groups
- [Moving Data Files Between Oracle ASM Disk Groups Using RMAN](#) (page 8-10) for information about moving database files with RMAN
- [Dropping Disk Groups](#) (page 4-69) for information about dropping a disk group
- *Oracle Clusterware Administration and Deployment Guide* for information about administering OCR and voting files
- *Oracle Database Administrator's Guide* for information about using the ALTER DATABASE MOVE DATAFILE SQL statement to move data files between Oracle ASM disk groups

4.8.4 Features Enabled By Disk Group Compatibility Attribute Settings

[Table 4-4](#) (page 4-58) describes the Oracle ASM features enabled by valid combinations of the disk group compatibility attribute settings.

The following list applies to [Table 4-4](#) (page 4-58) and [Table 4-5](#) (page 4-59).

- The value of COMPATIBLE.ASM must always be greater than or equal to the value of COMPATIBLE.RDBMS and COMPATIBLE.ADVM.
- A value of not applicable (n/a) means that the setting of the attribute has no effect on the feature.
- Oracle ASM and Oracle ACFS features not explicitly listed in [Table 4-4](#) (page 4-58) and [Table 4-5](#) (page 4-59) do not require advancing the disk group compatibility attribute settings.
- Oracle ASM and Oracle ACFS features explicitly identified by an operating system in [Table 4-4](#) (page 4-58) and [Table 4-5](#) (page 4-59) are available on that operating system starting with the associated disk group attribute settings.
- If encryption is configured for the first time on Oracle ASM 11g Release 2 (11.2.0.3) on Linux or if encryption parameters must be changed or a new volume encryption key must be created following a software upgrade to Oracle ASM 11g Release 2 (11.2.0.3) on Linux, then the disk group compatibility attributes for ASM and ADVM must be set to 11.2.0.3 or higher.
- All database files for Oracle Database can be stored in Oracle ACFS starting with Oracle Database 11g Release 2 (11.2.0.4), except for data files and redo logs in an Oracle Restart (standalone) configuration. Using replication or encryption with database files on Oracle ACFS is not supported.

Table 4-4 Oracle ASM features enabled by disk group compatibility attribute settings

Disk Group Features Enabled	COMPATIBLE.ASM	COMPATIBLE.RDBMS	COMPATIBLE.ADVM
Support for larger AU sizes (32 or 64 MB)	>= 11.1	>= 11.1	n/a
Attributes are displayed in the V \$ASM_ATTRIBUTE 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 files in a disk group	>= 11.2	n/a	n/a
Sector size set to nondefault 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
ASM_POWER_LIMIT value up to 1024	>= 11.2.0.2	n/a	n/a
Content type of a disk group	>= 11.2.0.3	n/a	n/a
Replication status of a disk group	>= 12.1	n/a	n/a
Managing a shared password file in a disk group	>= 12.1	n/a	n/a
Greater than 2 TB Oracle ASM disks without Oracle Exadata storage	>= 12.1	>= 12.1	n/a
Appliance mode support	>= 12.1.0.2	n/a	n/a
LOGICAL_SECTOR_SIZE	>= 12.2	n/a	n/a
Altering sector size	>= 12.2	n/a	n/a
Oracle ASM flex and extended disk groups	>= 12.2	>= 12.2	n/a
SCRUB_ASYNC_LIMIT	>= 12.2	n/a	n/a
PREFERRED_READ.ENABLED	>=12.2	n/a	n/a
Oracle Database 12c Release 2 (12.2) is required.			

Table 4-5 (page 4-59) describe the Oracle ACFS features enabled by valid combinations of the disk group compatibility attribute settings.

Table 4-5 Oracle ACFS features enabled by disk group compatibility attribute settings

Disk Group Features Enabled	COMPATIBLE.ASM	COMPATIBLE.RDBMS	COMPATIBLE.ADVM
Volumes in disk groups	>= 11.2	n/a	>= 11.2
Encryption, replication, security, tagging (Linux systems)	>= 11.2.0.2	n/a	>= 11.2.0.2
Read-only snapshots	>= 11.2.0.2	n/a	>= 11.2.0.2
Encryption, replication, security, tagging (Windows systems)	>= 11.2.0.3	n/a	>= 11.2.0.3
Read-write snapshots	>= 11.2.0.3	n/a	>= 11.2.0.3
Unlimited file system expansion	>= 11.2.0.4	n/a	>= 11.2.0.4
Performance and scalability improvements in <code>ls</code> and <code>find</code>	>= 11.2.0.4	n/a	>= 11.2.0.4
Storing database data files and redo logs in Oracle ACFS file systems (on Windows the software must be 12.1.0.2 or higher and <code>COMPATIBLE.ADVM</code> must be set to 12.1.0.2)	>= 12.1	n/a	>= 12.1
Encryption, replication, security, tagging (Solaris systems)	>= 12.1	n/a	>= 12.1
Replication and tagging (AIX systems)	>= 12.1	n/a	>= 12.1
Creation from an existing snapshot and conversion of a snapshot	>= 12.1	n/a	>= 12.1
Support for 1023 snapshots	>= 12.1	n/a	>= 12.1.0.2
Accelerator volume for Oracle ACFS file system	>= 12.2	n/a	>= 12.2
Metadata storage on accelerator volume for Oracle ACFS file system	>=12.2	n/a	>=12.2
Logical sector size of the Oracle ADVM volume	>= 12.2	n/a	>= 12.2
Oracle ACFS support for 4 K sectors	>=12.2	n/a	>=12.2
Oracle ACFS automatic resize	>=12.2, or >=12.1.0.2 on Oracle Data Appliance (ODA)	n/a	>=12.2, or >=12.1.0.2 on Oracle Data Appliance (ODA)
Oracle ACFS sparse files	>=12.2, or >=12.1.0.2 on Oracle Data Appliance (ODA)	n/a	>=12.2, or >=12.1.0.2 on Oracle Data Appliance (ODA)
Oracle ACFS compression	>=12.2	n/a	>=12.2
Cloning, duplicating, remastering, and setting quotas for snapshots snapshots	>= 12.2	n/a	>= 12.2

Table 4-5 (Cont.) Oracle ACFS features enabled by disk group compatibility attribute settings

Disk Group Features Enabled	COMPATIBLE.ASM	COMPATIBLE.RDBMS	COMPATIBLE.ADVM
Oracle ACFS snapshot-based replication	>=12.2	n/a	>=12.2
Oracle ACFS Defragger	>=12.2	n/a	>=12.2

See Also:

- [Overview of Oracle ACFS](#) (page 11-1) for information, including any limitations or restrictions, about storing data files in an Oracle ACFS file system
- [ASM_POWER_LIMIT](#) (page 3-11) for information about the rebalancing process and the ASM_POWER_LIMIT initialization parameter

4.8.5 Considerations When Setting Disk Group Compatibility

If you advance disk group compatibility, then you could enable the creation of files that are too large to be managed by a previous release of Oracle Database.

For example, 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 4-6](#) (page 4-60) shows the maximum Oracle ASM file sizes supported for multiple COMPATIBLE.RDBMS settings with the AU_SIZE disk group attribute set to 1 MB. This table 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 causes the file to be 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.

Table 4-6 Maximum Oracle ASM file sizes for disk groups with AU_SIZE set to 1 MB

Redundancy	COMPATIBLE.RDBMS = 10.1	COMPATIBLE.RDBMS >= 11.1
External	16 TB	128 TB
Normal	5.8 TB	93 TB
High	3.9 TB	62 TB

[Table 4-7](#) (page 4-60) shows the maximum Oracle ASM file sizes supported for multiple COMPATIBLE.RDBMS settings with the AU_SIZE disk group attribute set to 4 MB.

Table 4-7 Maximum Oracle ASM file sizes for disk groups with AU_SIZE set to 4 MB

Redundancy	COMPATIBLE.RDBMS = 10.1	COMPATIBLE.RDBMS >= 11.1
External	64 TB	128 TB

Table 4-7 (Cont.) Maximum Oracle ASM file sizes for disk groups with AU_SIZE set to 4 MB

Redundancy	COMPATIBLE.RDBMS = 10.1	COMPATIBLE.RDBMS >= 11.1
Normal	32 TB	128 TB
High	21 TB	128 TB

Note:

The 128 TB limit under the **COMPATIBLE.RDBMS >= 11.1** column in tables [Table 4-6](#) (page 4-60) and [Table 4-7](#) (page 4-60) is not a limit specific to Oracle ASM. The limit is specific to the database data file, which is calculated as $2^2 * \text{logical block size}$. The block size can be at most 32 KB, or 2^5 . The resulting calculation is $2^2 * 2^5 = 2^7$, or 128 TB.

See Also:

- [Oracle ASM Storage Limits](#) (page 4-49) for information about Oracle ASM storage sizes
- *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

4.9 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 must designate separate operating system group. You can manage Oracle ASM file access control with ASMCMD commands and SQL statements.

This section contains these topics:

- [About Oracle ASM File Access Control](#) (page 4-62)
- [Setting Disk Group Attributes for Oracle ASM File Access Control](#) (page 4-63)
- [Using SQL Statements to Manage Oracle ASM File Access Control](#) (page 4-64)

See Also:

- [Using Separate Operating System Groups for Oracle ASM Users](#) (page 3-29) for information about designating separate operating system groups
 - [ASMCMD File Access Control Commands](#) (page 10-90) for information about managing Oracle ASM File Access Control with ASMCMD commands
 - [Viewing Oracle ASM File Access Control Information](#) (page 6-5) for information about views that provide details about Oracle ASM file access control
 - [Authentication for Accessing Oracle ASM Instances](#) (page 3-27) for information about controlling accessing to Oracle ASM instances
-

4.9.1 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](#) (page 3-30) 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 can 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 you create a file, a group is not automatically associated with the file. You must use the `asmcmd chgrp` command to set the group manually because Oracle ASM File Access Control does not have the concept of the primary group of a user.

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.

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.

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.

3. Optionally, you can create user groups that are groups of database users that share the same access permissions to Oracle ASM files.

See Also:

- [Setting Disk Group Attributes for Oracle ASM File Access Control](#) (page 4-63) for information about setting disk group attributes
 - [Using SQL Statements to Manage Oracle ASM File Access Control](#) (page 4-64) for information about using `ALTER DISKGROUP SET PERMISSION` and `ALTER DISKGROUP SET OWNERSHIP` for managing Oracle ASM File Access Control with SQL statements
 - [Using SQL Statements to Manage Oracle ASM File Access Control](#) (page 4-64) for information about using `ALTER DISKGROUP ADD USERGROUP`
-
-

4.9.2 Setting Disk Group Attributes for Oracle ASM File Access Control

To manage Oracle ASM File Access Control for a disk group, you must set the `ACCESS_CONTROL.ENABLED` and `ACCESS_CONTROL.UMASK` disk group attributes. You can set the attributes by altering the disk group with the `ALTER DISKGROUP SQL` statement or you can set the attributes with the `ASMCMD setattr` command.

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 or the `ASMCMD chmod` command 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 4-11](#) (page 4-64) shows how to enable Oracle ASM File Access Control for a disk group using SQL*Plus. In this example, the `umask` permissions setting is `026` which enables read-write access for the owner (masks out nothing with 0), read access for users in the group (masks out write permission with 2), and no access to others (masks out all access with 6) not in the group.

Example 4-11 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';
```

See Also:

- [setattr](#) (page 10-75) for information about the ASMCMD `setattr` command
 - [chmod](#) (page 10-92) for information about the ASMCMD `chmod` command
 - [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility attributes
-
-

4.9.3 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.

When administering users with SQL*Plus, the users must be existing operating system users and their user names must have corresponding operating system user IDs. However, only users in the same cluster as the Oracle ASM instance can be validated.

Note:

When modifying ownership of files, the previous owner of a file cannot be dropped while the file is still open even after the ownership of the file has changed.

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 operating system users of 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. For 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. 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, for a database instance to read files in a disk group without creating any files, then you must use the ADD USER clause to add the operating system user associated with that database to the disk group. Also, you can use this clause to add the operating system user associated with a database 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 REPLACE USER`

Replaces one operating system user with another in an Oracle ASM disk group. The replacement user must not be a user currently in the disk group user list.

If the command succeeds, all files that were previously owned by current user are now owned by the replacement user. The current user is automatically removed from the user list of the disk group. Adding the current user back to the disk group is possible, but this user does not own any files.

When changing the owner of an open file, the new owner cannot be dropped with the owner's files until all instances within a cluster are restarted.

Only an Oracle ASM administrator can run this command.

```
SQL> ALTER DISKGROUP data REPLACE USER 'oracle1' WITH 'oracle2';
```

- `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.

```
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 run 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.

```
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*

4.10 Maintaining Disk Groups

This section discusses maintenance operations performed on disk groups.

This section contains the following topics:

- [Mounting and Dismounting Disk Groups](#) (page 4-67)
- [Mounting Disk Groups Using the FORCE Option](#) (page 4-67)
- [Checking the Internal Consistency of Disk Group Metadata](#) (page 4-68)
- [Dropping Disk Groups](#) (page 4-69)
- [Renaming Disks Groups](#) (page 4-70)

4.10.1 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 process makes the disk groups 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.

When you want to mount or dismount disk groups manually, use the `ALTER DISKGROUP . . . MOUNT` or `ALTER DISKGROUP . . . DISMOUNT` statement. You can mount or dismount disk groups by name, or specify `ALL`. You can also specify the optional `FORCE` option. For information about using the `FORCE` option, refer to "[Mounting Disk Groups Using the FORCE Option](#) (page 4-67)".

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;
```

4.10.2 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 cannot access are placed 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.

Note:

An Oracle ASM instance mounts an incomplete disk group differently depending on the specified compatibility.

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:

```
SQL> ALTER DISKGROUP data1 MOUNT FORCE;
```

See Also:

- [Oracle ASM Fast Mirror Resync](#) (page 4-43) for more information about setting the `DISK_REPAIR_TIME` disk group attribute
 - [Features Enabled By Disk Group Compatibility Attribute Settings](#) (page 4-57) for information features enabled by disk group compatibility attributes
 - *Oracle Database SQL Language Reference* for additional information about the `ALTER DISKGROUP` statement and the `FORCE` option
-
-

4.10.3 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

4.10.4 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.

If you cannot mount a disk group but must 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:

```
SQL> 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 disks from a disk group with `ASMCA`.

See Also:

- [Example 12-9](#) (page 12-5) for an example of a query on the `V $ASM_VOLUME` view
 - [Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#) (page 16-23) for information about deregistering and dismounting Oracle ACFS file systems
 - [Managing Disk Groups with ASMCA](#) (page 9-2) for information about ASMCA
-
-

4.10.5 Renaming Disk 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}] dgroupname=diskgroup
  newdgroupname=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 run. 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 rerun phase `two` using the generated configuration file.
- `dgname=diskgroup`

Specifies the name of the disk group to be renamed.
- `newdgname=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 run. 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 clean 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 are to 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 are to 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 behavior, 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) `srvctl remove diskgroup` command.

[Example 4-12](#) (page 4-72) shows several examples of the use of `renamedg`. The first example renames the `fra1` disk group to `fra2` 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`.

After renaming a disk group, you can rename the disks in the disk group to match the new disk group name. For example:

```
SQL> ALTER DISKGROUP fra2 RENAME DISKS ALL;
```

For information about renaming disks, refer to "[Renaming Disks in Disk Groups](#) (page 4-24)".

Example 4-12 Using renamedg

```
$ renamedg dname=fra1 newdname=fra2 asm_diskstring='/devices/disk*'
  verbose=true

$ renamedg phase=one dname=fra1 newdname=fra2
  asm_diskstring='/devices/disk*' config=/tmp/fra2.conf verbose=true

$ renamedg phase=two dname=fra1 newdname=fra2 config=/tmp/fra2.conf
  verbose=true
```

4.11 Managing Oracle ASM Flex Disk Groups

Oracle ASM provides database-oriented storage management with flex and extended disk groups.

An Oracle ASM flex disk group is a disk group type that supports Oracle ASM file groups. An Oracle ASM file group describes a group of files that belong to a database, and enables storage management to be performed at the file group, or database, level. An Oracle ASM extended disk group has all the features of a flex disk group, with a specific design to be highly available in an extended (or stretch) cluster environment.

This section discusses the following topics:

- [About Oracle ASM Flex Disk Groups](#) (page 4-73)
- [About Oracle ASM File Groups](#) (page 4-74)
- [About Oracle ASM File Group Properties](#) (page 4-76)
- [About Quota Groups for Oracle ASM File Groups](#) (page 4-79)
- [About Oracle ASM Extended Disk Groups](#) (page 4-80)

4.11.1 About Oracle ASM Flex Disk Groups

An Oracle ASM flex disk group is a disk group type that supports Oracle ASM file groups and quota groups.

In general, a flex disk group enables users to manage storage at the granularity of the database, in addition to at the disk group level.

The following are characteristics of a flex disk group:

- The redundancy of files in a flex disk group is flexible and enables storage management at the database level.
- File groups of flex disk groups are used to describe database files.

Each database has its own file group, and storage management can be done at the file group level, in addition to the disk group level. For example, you can specify different redundancy and rebalance settings for different file groups. File groups are associated with a quota group, enabling easy quota management.

- The redundancy setting of a flex disk group is set to `FLEX REDUNDANCY` and each file group of the flex disk group has its own redundancy property setting.
- A flex disk group generally tolerates two failures, the same as a `HIGH` redundancy disk group.

However, if the disk group has fewer than five failure groups, then it can only tolerate one failure. This restriction is the result of the quorum requirement of the disk group.

- The failure tolerated by a database is dependent on the redundancy setting of its associated file group in the flex disk group.
- A flex disk group requires a minimum of three failure groups.
- Migrating (converting) to a flex disk group can be performed for a `NORMAL` or `HIGH` redundancy disk group with the `CONVERT REDUNDANCY TO FLEX` clause of the `ALTER DISKGROUP SQL` statement. You cannot migrate an `EXTERNAL` redundancy disk group to a `FLEX` redundancy disk group.
- When migrating to a flex disk group, the disk group must have a minimum of three failure groups and must be mounted in restricted mode (`MOUNTED RESTRICTED`) for the duration of the migration.
- The default size of the allocation unit (AU) is 4 M.
- The Virtually Allocated Metadata (VAM) is enabled and required for a flex disk group.
- The `COMPATIBLE .ASM` and `COMPATIBLE .RDBMS` disk group attributes must be set to 12.2 or greater.

Note:

The `REQUIRED_MIRROR_FREE_MB` and `USABLE_FILE_MB` columns in the `V $ASM_DISKGROUP` view report no values for a flex disk group.

The following examples illustrate how to create and migrate to a flex disk group with SQL statements.

Example 4-13 Using CREATE DISKGROUP to create a flex disk group

```
SQL> CREATE DISKGROUP flex_data FLEX REDUNDANCY DISK my_disk_discovery_path;
```

Example 4-14 Using ALTER DISKGROUP to migrate a normal disk group to a flex disk group

```
SQL> ALTER DISKGROUP data MOUNT RESTRICTED;
```

```
SQL> ALTER DISKGROUP data CONVERT REDUNDANCY TO FLEX;
```

4.11.2 About Oracle ASM File Groups

An Oracle ASM file group is a group of files that share the same set of properties and characteristics.

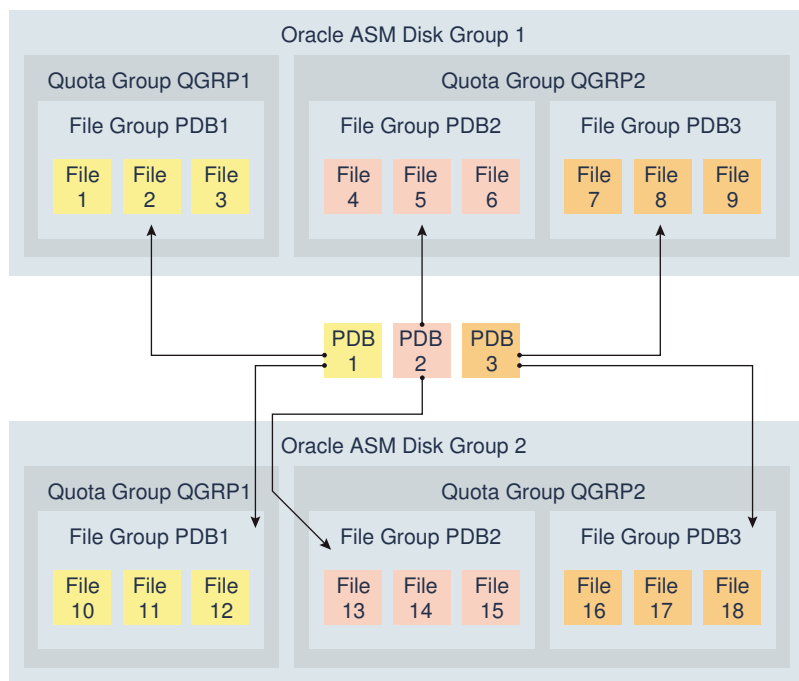
A major benefit of file groups is the ability to have a different availability specification for each database that shares the same disk group.

The set of properties for an Oracle ASM file group includes redundancy, rebalance rebuild priority, rebalance power limit, client compatibility, striping, quota group, and the access control list.

The following list provides important notes about file groups.

- A disk group contains at least one file group, the default file group.
- A disk group can contain multiple file groups.
- A disk group must have FLEX or EXTENDED redundancy to contain a file group.
- A disk group can store files belonging to multiple databases with each database having a separate file group.
- A database can have only one file group in a disk group.
- A database can span multiple disk groups with multiple file groups in different disk groups. The file groups belonging to a database in multiple disk groups should all have the same name for consistency and easy identification.
- A file group can belong to only one disk group.
- A file group can describe only one database, PDB, CDB, volume, or cluster.
- A file group can belong to only one quota group.
- Automatically created file groups are associated with the generic quota group.
- When a database, PDB, or CDB is created, if an existing file group has a client ID or name that matches the ID or name of the database, PDB, or CDB, then that file group is used to describe its files. Otherwise, a new file group is created for the database, PDB, or CDB.

The following figure illustrates the use of Oracle ASM file groups and quota groups in a multitenant environment. File groups and quota groups in different disk groups can have the same name.

Figure 4-1 Oracle ASM File Groups

In the figure, file groups named PDB1 in disk group 1 and disk group 2 are dedicated to pluggable database PDB1, file groups named PDB2 in disk group 1 and disk group 2 are dedicated to pluggable database PDB2, and file groups named PDB3 in disk group 1 and disk group 2 are dedicated to pluggable database PDB3.

File group PDB1 in disk group 1 belongs to quota group QGRP1 in disk group 1. File group PDB1 in disk group 2 belongs to quota group QGRP1 in disk group 2. File groups PDB2 and PDB3 in disk group 1 belong to quota group QGRP2 in disk group 1. File groups PDB2 and PDB3 in disk group 2 belong to quota group QGRP2 in disk group 2.

You can administer file groups with SQL statements. For example, you can use the `FILEGROUP` clause of the `ALTER DISKGROUP` SQL statement to administer file groups. [Example 4-15](#) (page 4-75) illustrates how SQL statements are used to add a file group for a pluggable database. The SQL statements also associate the newly added file group to an existing quota group. [Example 4-16](#) (page 4-75) illustrates how to modify the redundancy property for a specific file type in a file group.

Example 4-15 Adding a File Group to a Database

```
ALTER DISKGROUP DiskGroup_1 ADD FILEGROUP FileGroup_PDB1
    DATABASE PDB1
    SET 'quota_group' = 'QuotaGroup_QGRP1';
```

```
ALTER DISKGROUP DiskGroup_1 ADD FILEGROUP FileGroup_PDB2
    DATABASE PDB2
    SET 'quota_group' = 'QuotaGroup_QGRP2';
```

Example 4-16 Modifying a File Group

```
ALTER DISKGROUP DiskGroup_1 MODIFY FILEGROUP FileGroup_PDB1
    SET 'datafile.redundancy' = 'HIGH';
```

ASMCMD also supports commands to manage and track file and quota groups. These commands include:

- `chfg` and `chgg` to modify a file group or quota group
- `lsfg` and `lsqg` to list file groups or quota groups
- `mkfg` and `mkqg` to create a file group or quota group
- `rmfg` and `rmqg` to delete a file group or quota group
- `mvfg` to move a file group

To monitor Oracle ASM file groups, you can check the `V$ASM_FILEGROUP`, `V$ASM_FILEGROUP_PROPERTY`, and `V$ASM_FILEGROUP_FILE` views.

See Also:

- [About Quota Groups for Oracle ASM File Groups](#) (page 4-79) for information about quota groups.
 - [ASMCMD File Group Management Commands](#) (page 10-78) for information about ASMCMD commands to administer Oracle ASM File Groups
 - *Oracle Database Reference* for information about the `V$ASM*` dynamic performance views
-

4.11.3 About Oracle ASM File Group Properties

The following list summarizes the properties of file groups.

- `COMPATIBLE.CLIENT`

This is a file group property. The current value is the last known client compatibility during a file open operation. The default value for `COMPATIBLE.CLIENT` is `12.2.0.0.0`.

The `COMPATIBLE.RDBMS` disk group attribute must be set less than or equal to:

- `COMPATIBLE.CLIENT` of each file group (when set)
- `COMPATIBLE` database initialization parameter (for each database that is connected to Oracle ASM)
- `COMPATIBLE.ASM` of the disk group

You must be an Oracle ASM administrator to modify this property.

- `OWNER`

This is a file group property. To set the `OWNER` property, the `ACCESS_CONTROL.ENABLED` disk group attribute must be set to `TRUE` to enable Oracle ASM File Access Control. The user specified as the owner must be a member of the disk group.

Only the Oracle ASM administrator, the `OWNER`, the `OWNER` of the `USER_GROUP` associated with the file group, and a member of the `USER_GROUP` associated with the file group can modify the attributes of file groups that are labeled with `A` database administrator can modify this property. This is true only if `ACCESS_CONTROL.ENABLED` is true.

You can display the owner with the `USER_NUMBER` and `USER_INCARNATION` columns of the `V$ASM_FILEGROUP` view. The values in `USER_NUMBER` and `USER_INCARNATION` can be joined with the `USER_NUMBER` and `INCARNATION` columns of the `V$ASM_USER` view to display the user names.

The default value is `NULL`.

You must be an Oracle ASM administrator to modify this property.

- `POWER_LIMIT`

This is a file group property. This property specifies the rebalance power, which determines the number of parallel relocations for disk rebalancing of this file group. The valid settings are 1 to 1024. When rebalancing a file group, the power limit value is applied as follows:

1. If the power option has been explicitly specified in the `ALTER DISKGROUP REBALANCE SQL` statement or the `ASMCMD rebal` command, then the power value specified in the SQL rebalance statement or `ASMCMD rebalance` command is used for rebalancing all file groups.
2. If the power option has not been specified in the SQL rebalance statement or `ASMCMD rebalance` command, then the file group `POWER_LIMIT` property value, if set, applies for the file group rebalance.
3. If the power option has not been specified in the SQL rebalance statement or `ASMCMD rebalance` command, and the `POWER_LIMIT` property has not been set for the file group, then the power value defaults to the value of the `ASM_POWER_LIMIT` initialization parameter for the file group rebalance.

You must be a database administrator to modify this property.

- `PRIORITY`

This is a file group property. This property specifies the rebalance priority and ensures that certain types of files are rebalanced before others. File groups with the highest priority are completely rebalanced first, then file groups with the next highest priority, and so on to the lowest priority.

You can set the rebalance priority setting to `HIGHEST`, `HIGH`, `MEDIUM`, `LOW`, or `LOWEST`. The default value is `MEDIUM`.

You must be a database administrator to modify this property.

- `QUOTA_GROUP`

This is a file group property. This property specifies the name of the quota group to which this file group belongs. The valid setting is an existing quota group name. The default value is `GENERIC`. Every file group in the disk group must be a member of a quota group.

You must be an Oracle ASM administrator to modify this property.

- `REDUNDANCY`

This is a file type property. The default redundancy setting for each file type is derived from the system templates. The valid settings are `HIGH`, `MIRROR`, and `UNPROTECTED`.

The redundancy of a file can be changed after the file is created. If redundancy has been changed, then the `REMIRROR` column of `V$ASM_FILE` contains `Y` to indicate that the file needs new mirroring which initiates a rebalance is run to put the new

redundancy into effect. After the rebalance completes, the value in the `REMIRROR` column contains `N`.

If the redundancy in the file group is increased, then the value of the `REDUNDANCY` column of `V$ASM_FILE` is increased after rebalance completes. However, if the redundancy in the file group is decreased, then the value of the `REDUNDANCY` column of `V$ASM_FILE` is lowered immediately before rebalance is run.

Changing the redundancy on a file type granularity or on the entire file group affects all file types.

You must be a database administrator to modify this property.

- **STRIPING**

This is a file type property, and is set for each file type. Usually the default value for each file type is sufficient and is not changed.

Most files have a `COARSE` stripe, except for the control file which has a `FINE` stripe by default. The valid settings are `COARSE` and `FINE`. The default value is `COARSE` or derived from the file type template.

`STRIPING` affects only newly-created files and cannot be changed after the file is created.

You must be a database administrator to set this property.

- **USER_GROUP**

This is a file group property. To set the `USER_GROUP` property, the `ACCESS_CONTROL.ENABLED` disk group attribute must be set to `TRUE` to enable Oracle ASM File Access Control. The user group specified must be an existing user group.

Only the Oracle ASM administrator, the `OWNER`, the `OWNER` of the `USER_GROUP` associated with the file group, and a member of the `USER_GROUP` associated with the file group can modify the attributes of file groups that are labeled with `A`. A database administrator can modify this property. This is true only if `ACCESS_CONTROL.ENABLED` is true.

You can display the user group information with the `USERGROUP_NUMBER` and `USERGROUP_INCARNATION` columns of the `V$ASM_FILEGROUP` view. The values in `USERGROUP_NUMBER` and `USERGROUP_INCARNATION` can be joined with the `USERGROUP_NUMBER` and `INCARNATION` columns of the `V$ASM_USERGROUP` view to display the user group names.

Files created in a file group inherit the `USER_GROUP` property setting.

The default value is `NULL`.

You must be an Oracle ASM administrator to modify this property.

The following example illustrates how to set the `REDUNDANCY` file group property for the datafile and archivelog file types.

Example 4-17 Setting a file group property

```
ALTER DISKGROUP DiskGroup_2 MODIFY FILEGROUP FileGroup_PDB1
SET 'datafile.redundancy' = 'HIGH';
```

```
ALTER DISKGROUP DiskGroup_2 MODIFY FILEGROUP FileGroup_PDB1
SET 'archivelog.redundancy' = 'MIRROR';
```

See Also:

- [What Types of Files Does Oracle ASM Support?](#) (page 5-1) for a list of Oracle ASM file types
 - [Managing Oracle ASM File Access Control for Disk Groups](#) (page 4-61) for information about Oracle ASM File Access Control
-

4.11.4 About Quota Groups for Oracle ASM File Groups

A quota group defines the quota allocated to a group of Oracle ASM File Groups. A file group belongs to only one quota group. For information about file groups, refer to [About Oracle ASM File Groups](#) (page 4-74).

A quota can describe an aggregate of space used by different databases.

The following list provides important notes about quota groups for file groups.

- A file group can belong to only one quota group.
- A quota group cannot span multiple disk groups.
- A quota group describes an aggregate of space used by one file group or multiple file groups in the same disk group.

A quota is enforced when a file is created and resized.

- A quota is a physical space.

As a consequence, if the quota limit has been set to 10 M, then a two-way mirrored 6 M file exceeds the quota.

- Each quota group has two values: the limit and the current used space.

The limit value can be set below the current used space. This action prevents any additional space being allocated for files described by file groups associated with this quota group.

- A file group can be moved from one quota group to another, regardless whether or not the target quota group has enough space for the file group.

This capability enables an Oracle ASM administrator to stop any files described by a specific file group from allocating additional space.

You can administer quota groups with SQL and ASMCMD commands. The following examples illustrate how SQL statements are used to add a quota group to a disk group and how to move a file group from its current quota group to a new quota group.

Example 4-18 Adding a Quota Group to a Disk Group

```
ALTER DISKGROUP DiskGroup_2 ADD QUOTAGROUP QuotaGroup_QGRP3
SET 'quota' = 10m;
```

Example 4-19 Moving a File Group to a new Quota Group

```
ALTER DISKGROUP DiskGroup_2 MODIFY FILEGROUP FileGroup_PDB1
SET 'quota_group' = 'QuotaGroup_QGRP3';
```

4.11.5 About Oracle ASM Extended Disk Groups

An Oracle ASM extended disk group has all the features of a flex disk group, with a specific design to be highly available in an extended (stretch) cluster environment, which contains nodes that span multiple physically separated sites.

The following are characteristics of an extended disk group:

- The redundancy setting of an extended disk group is set to `EXTENDED REDUNDANCY` and each file group of the extended disk group has its own redundancy property setting.
- The redundancy of files and file groups in an extended disk group is specified for each site, rather than for each disk group.

This means that the number of copies determined by the `REDUNDANCY` setting is multiplied by the number of data sites. If the `REDUNDANCY` is set to `MIRROR`, then that specifies two copies, resulting in four copies total for two data sites.
- An extended disk group can tolerate the loss of an entire site, plus the loss of up to two failure groups in another site.

This means that an extended disk group can tolerate failures at the site level in addition to the failure group level.
- All data sites must have the same number of failure groups during disk group creation.
- A quota group is specified for each disk group, rather than for each site.
- The quota group limit is the physical space required for all copies across all sites.

For example, with two data sites, a 6 M file with redundancy set to `MIRROR` uses 24 M of the quota limit.
- There must be a total of three sites, two data sites and one quorum site, to create an extended disk group. Each data site should have three failure groups, and the quorum site should have one failure group.
- For database files, the redundancy setting is dictated by the redundancy setting of the file groups of the extended disk group.
- The `COMPATIBLE . ASM` and `COMPATIBLE . RDBMS` disk group attributes must be set to 12 . 2 or greater.
- The minimum allocation unit (AU) size is 4 M.
- Converting from another disk group type to an extended disk group is not supported.
- The `SITE` specification for failure groups can be inferred from cells if Oracle Exadata has been configured. Otherwise, the site specification can be specified with the SQL statement when creating a disk group.

Note:

The `REQUIRED_MIRROR_FREE_MB` and `USABLE_FILE_MB` columns in the `V $ASM_DISKGROUP` view report no values for an extended disk group.

The following example illustrates how to create an extended disk group with a SQL statement.

Example 4-20 Using CREATE DISKGROUP with extended disk groups

```
SQL> CREATE DISKGROUP extended_site_data EXTENDED REDUNDANCY
      SITE NY FAILGROUP fg1 DISK '/devices/disks/disk01'
           FAILGROUP fg2 DISK '/devices/disks/disk02'
           FAILGROUP fg3 DISK '/devices/disks/disk03'
      SITE NJ FAILGROUP fg4 DISK '/devices/disks/disk04'
           FAILGROUP fg5 DISK '/devices/disks/disk05'
           FAILGROUP fg6 DISK '/devices/disks/disk06'
      SITE QM QUORUM
           FAILGROUP fg7 DISK '/devices/disks/disk07';
```

Administering Oracle ASM Files, Directories, and Templates

Administration of Oracle ASM includes the management of 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?](#) (page 5-1)
- [About Oracle ASM File Names](#) (page 5-3)
- [Creating and Referencing Oracle ASM Files in the Database](#) (page 5-9)
- [Managing Alias Names for Oracle ASM File Names](#) (page 5-14)
- [Managing Disk Group Directories](#) (page 5-16)
- [Accessing Oracle ASM Files with the XML DB Virtual Folder](#) (page 5-17)
- [Using DBMS_FILE Transfer Utility for Oracle ASM](#) (page 5-19)
- [Managing Disk Group Templates](#) (page 5-19)

5.1 What Types of Files Does Oracle ASM Support?

Oracle ASM supports most file types required by the database.

[Table 5-1](#) (page 5-1) lists file types that are supported directly by Oracle ASM and lists the system default template that provides the attributes for file creation. Some file types shown in the table are related to specific products or features, and are not discussed in this book.

Table 5-1 *File types supported by Oracle ASM*

File Type	Default Templates
Archive log files	ARCHIVELOG
Oracle ASM Persistent initialization parameter file (SPFILE)	ASMPARAMETERFILE
Oracle ASM Persistent initialization parameter file (SPFILE) backup	ASMPARAMETERFILEBACKUP
Audit spill files	AUDIT_SPILLFILES
Automatically generated control file backup	AUTOBACKUP
Auto-login software keystore	AUTOLOGIN_KEY_STORE

Table 5-1 (Cont.) File types supported by Oracle ASM

File Type	Default Templates
Archive log backup piece	BACKUPSET
Data file backup pieces	BACKUPSET
Data file incremental backup pieces	BACKUPSET
Change tracking file	CHANGETRACKING
Control files	CONTROLFILE
Data file copy	DATAFILE
Data files	DATAFILE
Data Guard configuration file	DATAGUARDCONFIG
Data Pump dumpset	DUMPSET
Flashback logs	FLASHBACK
Flash file	FLASHFILE
Cross-platform transportable data files incremental backup	INCR XTRANSPORT BACKUPSET
Software keystore	KEY_STORE
Oracle Cluster Registry file	OCRFILE
Redo log files	ONLINELOG
Persistent initialization parameter file (SPFILE)	PARAMETERFILE
Temporary files	TEMPFILE
Voting files	VOTINGFILE
Cross-platform transportable data files	XTRANSPORT
Cross-platform transportable data files backup	XTRANSPORT BACKUPSET
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 data files, 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 Also:

- [Introducing Oracle ACFS and Oracle ADVM](#) (page 11-1) for information files supported by Oracle ACFS
 - [Managing Disk Group Templates](#) (page 5-19) for a description of the system default templates
 - [Example 6-10](#) (page 6-6) for an example of the output of the `V $ASM_TEMPLATE` view
 - [lstmpl](#) (page 10-88) for an example of the output of the `ASMCMD lstmpl` command
-

5.2 About Oracle ASM File Names

Every file created in Oracle ASM is provided a file name generated by Oracle Managed Files, known as a fully qualified file name (or system-generated file name). The fully qualified file name represents a complete path name in the Oracle ASM file system. An example of a fully qualified file name is:

```
+data/orcl/controlfile/Current.256.541956473
```

You can use the fully qualified file name to reference (read or retrieve) an Oracle ASM file. You can also use other abbreviated file name formats, such as an alias Oracle ASM file name, to reference an Oracle ASM file.

Oracle ASM generates a fully qualified file name upon any request to create a file. A creation request does not or cannot specify a fully qualified file name. 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 file name. If you specify an alias in the creation request, Oracle ASM also creates the alias so that it references the fully qualified file name.

Oracle ASM file creation requests are either single file creation requests or multiple file creation request.

Note:

- 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 must reference an Oracle ASM file in a SQL statement. As with other Oracle Database file names, Oracle ASM file names are kept in the control file and the RMAN catalog.
 - You can use fully qualified and numeric file names in single-file creation if you specify the `REUSE` keyword, as described in [Using Oracle ASM File Names in SQL Statements](#) (page 5-10).
 - Forward slash (/) and backslash (\) are interchangeable in file names. File names are case insensitive, but case retentive.
-

[Table 5-2](#) (page 5-4) specifies the valid usage for each file name form. If the form is used for file creation, the table specifies whether the created file is an Oracle Managed Files (OMF).

Table 5-2 Valid usage for Oracle ASM file name forms

File Name 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

This section contains the following topics:

- [Single File Creation Form](#) (page 5-4)
- [Multiple File Creation Form](#) (page 5-4)
- [Fully Qualified File Name Form](#) (page 5-5)
- [Alias Oracle ASM File Name Forms](#) (page 5-6)
- [Creating a Tablespace in Oracle ASM: Using a Data File with an Alias Name](#) (page 5-7)
- [Alias Oracle ASM File Name with Template Form](#) (page 5-7)
- [Incomplete Oracle ASM File Name Form](#) (page 5-8)
- [Incomplete Oracle ASM File Name with Template Form](#) (page 5-8)

5.2.1 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 file name 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 file name is called for in a SQL statement, such as `CREATE TABLESPACE`, `CREATE CONTROLFILE`.

5.2.2 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 run a `CREATE TABLESPACE` statement (without a file name 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 file name 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 file name (for example, `+data`), and whenever a command is run 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 file name in other *_DEST initialization parameters.

5.2.3 Fully Qualified File Name Form

You can use a fully qualified form 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 5-3](#) (page 5-5).
- *filetypetag* is type-specific information about the file and can be one of the tags shown in [Table 5-3](#) (page 5-5).
- *file.incarnation* is the file/incarnation pair, used to ensure uniqueness.

An example of a fully qualified Oracle ASM file name is:

```
+data/orcl/controlfile/Current.256.541956473
```

Table 5-3 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 Backup	Archive 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>	
TEMPFILE	Tempfiles	<i>tsname</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.

Table 5-3 (Cont.) Oracle file types and Oracle ASM file type tags

Oracle ASM File Type	Description	Oracle ASM File Type Tag	Comments
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 where 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

5.2.4 Alias Oracle ASM File Name Forms

You can use Alias Oracle ASM file names 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 file names 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 file name, the file is created with a fully qualified name, and the alias file name is additionally created. You can then access the file with either name.

Alias Oracle ASM file names are distinguished from fully qualified file names or numeric file names 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 file names are:


```
+data/orcl/control_file1
+fra/recover/second.dbf
```

Oracle Database references database files by their alias file names, 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 file names. The following are examples of how the database uses alias file names:

- Alias file names appear in V\$ views. For example, if you create a tablespace and use an alias file name for the data file, the V\$DATAFILE view shows the alias file name.
- When a control file points to data files and online redo log files, it can use alias file names.
- The CONTROL_FILES initialization parameter can use the alias file names of the control files. Database Configuration Assistant (DBCA) creates control files with alias file names.

Note:

Files created using an alias file name 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 file names, see "[Managing Alias Names for Oracle ASM File Names](#) (page 5-14)".

5.2.5 Creating a Tablespace in Oracle ASM: Using a Data File with an Alias Name

[Example 5-1](#) (page 5-7) 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.

Because an alias is used when creating the data file, the file is not an Oracle Managed Files (OMF) file and the file is not 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';
```

Example 5-1 *Creating a tablespace using an alias name*

```
CREATE UNDO TABLESPACE myundo
  DATAFILE '+data(my_undo_template)/orcl/my_undo_ts' SIZE 200M;
```

5.2.6 Alias Oracle ASM File Name with Template Form

An alias Oracle ASM file name with template is used only for Oracle ASM file creation operations. It has the following format:

```
+diskgroup(template name)/alias
```

Alias file names with template behave identically to alias file names except that a file created with an alias file name 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.

An example of an alias Oracle ASM file name 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 file name with template are not considered Oracle Managed Files and might require manual deletion in the future if they are no longer needed.

See Also:

- [Example 5-1](#) (page 5-7) for information about creating objects using an alias name
 - [Managing Disk Group Templates](#) (page 5-19) for information about the creation and maintenance of Oracle ASM templates
-

5.2.7 Incomplete Oracle ASM File Name Form

Incomplete Oracle ASM file names 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, then the DATAFILE template is used.

An example of using an incomplete Oracle ASM file name 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](#) (page 5-9)" for more information.

5.2.8 Incomplete Oracle ASM File Name with Template Form

Incomplete Oracle ASM file names 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 file name with template is setting the `DB_CREATE_FILE_DEST` initialization parameter to:

```
+data(mytemplate)
```

5.3 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](#) (page 5-9)
- [Using Oracle ASM File Names in SQL Statements](#) (page 5-10)
- [Moving Data Files Between Disk Groups With ALTER DATABASE](#) (page 5-11)
- [Oracle Database File Mapping for Oracle ASM Files](#) (page 5-11)

5.3.1 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 file name, 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 file names 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 file names 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
-
-

5.3.2 Using Oracle ASM File Names in SQL Statements

You can specify Oracle ASM file names in the file specification clause of your SQL statements. If you are creating a file for the first time, then use the creation form of an Oracle ASM file name. If the Oracle ASM file exists, then you must use the reference form of the file name, 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](#) (page 21-3)".

If a reference form is used with the REUSE keyword and the file does not exist, an error results.

[Example 5-2](#) (page 5-10) is an example of specifying an Oracle ASM file name in a SQL statement. In this case, it is used in the file creation form.

The tablespace mytblspace is created and comprises 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. You can use an AUTOEXTEND clause to override this default.

For more examples of the use of Oracle ASM file name in SQL statements, see [Creating Database Files in Oracle ASM Using SQL*Plus](#) (page 21-1).

Example 5-2 Using an Oracle ASM File name in a SQL statement

```
CREATE TABLESPACE mytblspace DATAFILE '+data' SIZE 200M AUTOEXTEND ON;
```

5.3.3 Moving Data Files Between Disk Groups With ALTER DATABASE

You can use the `ALTER DATABASE MOVE DATAFILE` SQL statement to move data files between disk groups while the database is open and users are accessing the data files.

[Example 5-3](#) (page 5-11) shows how you can move a data file from the DATA disk group to the DATA2 disk group.

Example 5-3 Moving data files online with ALTER DATABASE MOVE DATAFILE

```
/* Create a new disk group DATA2 using ASMCA, ASMCMD, or SQL */
/* Then create appropriate directories in the DATA2 disk group */

ALTER DISKGROUP data2 ADD DIRECTORY '+DATA2/ORCL';

ALTER DISKGROUP data2 ADD DIRECTORY '+DATA2/ORCL/DATAFILE';

/* Move the EXAMPLE data file in DATA to EXAMPLE_STORAGE in DATA2
ALTER DATABASE MOVE DATAFILE '+DATA/ORCL/DATAFILE/EXAMPLE.266.798707687'
    TO '+DATA2/ORCL/DATAFILE/EXAMPLE_STORAGE';

SELECT SUBSTR(name,1,64) FROM V$DATAFILE;

SUBSTR(NAME,1,64)
-----
+DATA/ORCL/DATAFILE/system.258.798707555
+DATA2/ORCL/DATAFILE/example_storage
+DATA/ORCL/DATAFILE/sysaux.257.798707491
+DATA/ORCL/DATAFILE/undotbs1.260.798707621
+DATA/ORCL/DATAFILE/users.259.798707621
```

See Also:

- *Oracle Database Administrator's Guide* for more information about using the `ALTER DATABASE MOVE DATAFILE` SQL statement to move data files
 - *Oracle Database SQL Language Reference* for information about using the `ALTER DATABASE` SQL statement.
 - [Moving Data Files Between Oracle ASM Disk Groups Using RMAN](#) (page 8-10) for information about using RMAN to move data files between disk groups
-

5.3.4 Oracle Database File Mapping for Oracle ASM Files

For an understanding of I/O performance, you must have detailed knowledge of the storage hierarchy in which files reside.

Oracle Database provides a set of dynamic performance views to show a complete mapping of a file to intermediate layers of logical volumes to actual physical devices. Using these dynamic views, you can locate the exact disk on which any block of a file resides.

Oracle Database communicates with a background process called FMON to manage the mapping information. Oracle provides the PL/SQL `DBMS_STORAGE_MAP` package to invoke mapping operations that populate the mapping views.

Oracle Database file mapping does not require third party libraries when mapping Oracle ASM files. In addition, Oracle Database supports Oracle ASM file mapping on all operating system platforms.

All the example in this section are run on the Oracle Database instance.

, refer to "".

This section contains the following topics:

- [Enabling File Mapping With Oracle ASM Files](#) (page 5-12)
- [Using File Mapping Views With Oracle ASM Files](#) (page 5-13)
- [Using the DBMS_STORAGE_MAP Package With Oracle ASM Files](#) (page 5-13)

See Also:

- [Oracle ACFS Support for Oracle Database File Mapping Views](#) (page 12-6) for information about file mapping with Oracle ACFS files
 - *Oracle Database Administrator's Guide* for additional information about Oracle Database file mapping
-

5.3.4.1 Enabling File Mapping With Oracle ASM Files

To enable file mapping, set the `FILE_MAPPING` initialization parameter to `TRUE`. The database instance does not have to be shut down to set this parameter. You can set the parameter using the following `ALTER SYSTEM` statement:

```
SQL> ALTER SYSTEM SET FILE_MAPPING=TRUE;
```

Run the appropriate `DBMS_STORAGE_MAP` mapping procedure.

- In a cold startup scenario, the Oracle Database has just started and no mapping operation has been invoked yet. You can execute the `DBMS_STORAGE_MAP.MAP_ALL` procedure to build the mapping information for the entire I/O subsystem associated with the database.

For example, the following command builds mapping information and provides for 10000 extents:

```
SQL> EXECUTE DBMS_STORAGE_MAP.MAP_ALL(10000);
```

- In a warm start scenario where the mapping information has already been built, you have the option to invoke the `DBMS_STORAGE_MAP.MAP_SAVE` procedure to save the mapping information in the data dictionary.

This procedure is invoked in `DBMS_STORAGE_MAP.MAP_ALL()` by default. This forces all of the mapping information in the SGA to be flushed to disk. The `DBMS_STORAGE_MAP.MAP_SAVE` procedure is invoked in `DBMS_STORAGE_MAP.MAP_ALL()` by default.

After you restart the database, use `DBMS_STORAGE_MAP.RESTORE()` to restore the mapping information into the SGA. If needed, `DBMS_STORAGE_MAP.MAP_ALL()` can be called to refresh the mapping information.

5.3.4.2 Using File Mapping Views With Oracle ASM Files

Mapping information generated by DBMS_STORAGE_MAP package is captured in dynamic performance views. These views include V\$MAP_COMP_LIST, V\$MAP_ELEMENT, V\$MAP_EXT_ELEMENT, V\$MAP_FILE, V\$MAP_FILE_EXTENT, V\$MAP_FILE_IO_STACK, V\$MAP_LIBRARY, and V\$MAP_SUBELEMENT.

For example, you can view file mappings with V\$MAP_FILE.

```
SQL> SELECT FILE_MAP_IDX, SUBSTR(FILE_NAME,1,45), FILE_TYPE, FILE_STRUCTURE FROM V$MAP_FILE;
```

FILE_MAP_IDX	SUBSTR(FILE_NAME,1,45)	FILE_TYPE	FILE_STRU
0	+DATA/ORCL/DATAFILE/system.258.841314095	DATAFILE	ASMFILE
1	+DATA/ORCL/DATAFILE/sysaux.257.841314021	DATAFILE	ASMFILE
2	+DATA/ORCL/DATAFILE/undotbs1.260.841314161	DATAFILE	ASMFILE
3	+DATA/ORCL/DATAFILE/example.266.841314219	DATAFILE	ASMFILE
4	+DATA/ORCL/DATAFILE/users.259.841314161	DATAFILE	ASMFILE
5	+DATA/ORCL/ONLINELOG/group_3.264.841314209	LOGFILE	ASMFILE
6	+FRA/ORCL/ONLINELOG/group_3.259.841314211	LOGFILE	ASMFILE
7	+DATA/ORCL/ONLINELOG/group_2.263.841314207	LOGFILE	ASMFILE
8	+FRA/ORCL/ONLINELOG/group_2.258.841314207	LOGFILE	ASMFILE
9	+DATA/ORCL/ONLINELOG/group_1.262.841314203	LOGFILE	ASMFILE
10	+FRA/ORCL/ONLINELOG/group_1.257.841314203	LOGFILE	ASMFILE
11	+DATA/ORCL/TEMPFILE/temp.265.841314217	TEMPFILE	ASMFILE
12	+DATA/ORCL/CONTROLFILE/current.261.841314199	CONTROLFILE	ASMFILE
13	+FRA/ORCL/CONTROLFILE/current.256.841314201	CONTROLFILE	ASMFILE

See Also:

Oracle Database Reference for information about dynamic views

5.3.4.3 Using the DBMS_STORAGE_MAP Package With Oracle ASM Files

You can use the procedures in the DBMS_STORAGE_MAP PL/SQL package to control mapping operations. For example, you can use the DBMS_STORAGE_MAP.MAP_OBJECT procedure to build mapping information for the database object that is specified by object name, owner, and type. After the DBMS_STORAGE_MAP.MAP_OBJECT procedure is run, then you can create a query to display the mapping information contained in the MAP_OBJECT view.

```
SQL> EXECUTE DBMS_STORAGE_MAP.MAP_OBJECT('EMPLOYEES','HR','TABLE');
```

```
SQL> SELECT io.OBJECT_NAME o_name, io.OBJECT_OWNER o_owner, io.OBJECT_TYPE o_type,
         mf.FILE_NAME, me.ELEM_NAME, io.DEPATH,
         (SUM(io.CU_SIZE * (io.NUM_CU - DECODE(io.PARITY_PERIOD, 0, 0,
         TRUNC(io.NUM_CU / io.PARITY_PERIOD)))) / 2) o_size
FROM MAP_OBJECT io, V$MAP_ELEMENT me, V$MAP_FILE mf
WHERE io.OBJECT_NAME = 'EMPLOYEES'
AND io.OBJECT_OWNER = 'HR'
AND io.OBJECT_TYPE = 'TABLE'
AND me.ELEM_IDX = io.ELEM_IDX
AND mf.FILE_MAP_IDX = io.FILE_MAP_IDX
GROUP BY io.ELEM_IDX, io.FILE_MAP_IDX, me.ELEM_NAME, mf.FILE_NAME, io.DEPATH,
         io.OBJECT_NAME, io.OBJECT_OWNER, io.OBJECT_TYPE
ORDER BY io.DEPATH;
```

O_NAME	O_OWNER	O_TYPE	FILE_NAME	ELEM_NAME	DEPTH	O_SIZE
--------	---------	--------	-----------	-----------	-------	--------

EMPLOYEES HR	TABLE	+DATA/ORCL/DATAFILE/example.266.841314219	+/devices/diskd2	0	64
EMPLOYEES HR	TABLE	+DATA/ORCL/DATAFILE/example.266.841314219	+/devices/diske2	0	64

See Also:

Oracle Database PL/SQL Packages and Types Reference for information about the DBMS_STORAGE_MAP package

5.4 Managing Alias Names for Oracle ASM File Names

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 file names.

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 (+).

This section contains the following sections:

- [Adding an Alias Name for an Oracle ASM File Name](#) (page 5-14)
- [Renaming an Alias Name for an Oracle ASM File Name](#) (page 5-14)
- [Dropping an Alias Name for an Oracle ASM File Name](#) (page 5-15)
- [Dropping Files and Associated Aliases from a Disk Group](#) (page 5-15)

See Also:

- *Oracle Database SQL Language Reference* for information about the ALTER DISKGROUP statement
 - [mkalias](#) (page 10-50) for information about creating aliases with ASMCMD
-

5.4.1 Adding an Alias Name for an Oracle ASM File Name

Use the ADD ALIAS clause of the ALTER DISKGROUP statement to create an alias name for an Oracle ASM file name. The alias name must consist of the full directory path and the alias itself.

[Example 5-4](#) (page 5-14) adds a new alias name for a system-generated file name.

Example 5-4 Adding an alias name for an Oracle ASM file name

```
ALTER DISKGROUP data ADD ALIAS '+data/orcl/second.dbf'
FOR '+data/orcl/datafile/mytable.342.123456789';
```

5.4.2 Renaming an Alias Name for an Oracle ASM File Name

Use the RENAME ALIAS clause of the ALTER DISKGROUP statement to rename an alias for an Oracle ASM file name. The old and the new alias names must consist of the full directory paths of the alias names.

[Example 5-5](#) (page 5-15) renames an alias.

Renaming an alias does not update the database reference to the alias.

Example 5-5 Renaming an alias name for an Oracle ASM file name

```
ALTER DISKGROUP data RENAME ALIAS '+data/orcl/datafile.dbf'
    TO '+data/payroll/compensation.dbf';
```

5.4.3 Dropping an Alias Name for an Oracle ASM File Name

Use the `DROP ALIAS` clause of the `ALTER DISKGROUP` statement to drop an alias for an Oracle ASM file name. 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 5-6](#) (page 5-15) drops an alias.

[Example 5-7](#) (page 5-15) fails because it attempts to drop a system-generated file name and this syntax is not allowed.

Example 5-6 Dropping an alias name for an Oracle ASM file name

```
ALTER DISKGROUP data DROP ALIAS '+data/payroll/compensation.dbf';
```

Example 5-7 Dropping an alias name for an Oracle ASM file name

```
ALTER DISKGROUP data
    DROP ALIAS '+data/orcl/datafile/mytable.342.123456789';
```

5.4.4 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 file name, or an alias name when specifying the file to delete. The underlying file on the file system is not dropped when the alias is dropped.

Some reasons why you may want 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 5-8](#) (page 5-15), the alias name for the file is used when deleting both the file and the alias from a disk group.

In [Example 5-9](#) (page 5-15), the Oracle Managed Files (system-generated) file name is used when dropping the file and any associated alias.

Example 5-8 Dropping a file and alias from a disk group using the alias name

```
ALTER DISKGROUP data DROP FILE '+data/payroll/compensation.dbf';
```

Example 5-9 Dropping a file and alias from a disk group using the system-generated name

```
ALTER DISKGROUP data
    DROP FILE '+data/orcl/datafile/mytable.342.123456789';
```

5.5 Managing Disk Group Directories

Oracle ASM disk groups contain a system-generated hierarchical directory structure for storing Oracle ASM files. The system-generated file name that Oracle ASM assigns to each file represents a path in this directory hierarchy. The following is an example of a system-generated file name:

```
+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 Directory](#) (page 5-16)
- [Renaming a Directory](#) (page 5-16)
- [Dropping a Directory](#) (page 5-17)

5.5.1 Creating a 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 5-10](#) (page 5-16) creates a hierarchical directory for disk group `data`, which can contain, for example, the alias name `+data/orcl/control_file1`.

Assuming no subdirectory exists under the directory `+data/orcl`, the SQL statement fails in [Example 5-11](#) (page 5-16).

Example 5-10 Creating a new directory

```
ALTER DISKGROUP data ADD DIRECTORY '+data/orcl';
```

Example 5-11 Creating a new subdirectory

```
ALTER DISKGROUP data
  ADD DIRECTORY '+data/orcl/first_dir/second_dir';
```

5.5.2 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 5-12](#) (page 5-17) renames a directory.

Example 5-12 Renaming a directory

```
ALTER DISKGROUP data RENAME DIRECTORY '+data/mydir'
      TO '+data/yourdir';
```

5.5.3 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 5-13](#) (page 5-17) deletes a directory along with its contents.

Example 5-13 Dropping a directory

```
ALTER DISKGROUP data DROP DIRECTORY '+data/yourdir' FORCE;
```

5.6 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. For example, `/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.

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:

- [cp](#) (page 10-44) for information about the `ASMCMD cp` command.
- [Inside /sys/asm](#) (page 5-18) for additional information about the `/sys/asm` folder
- *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

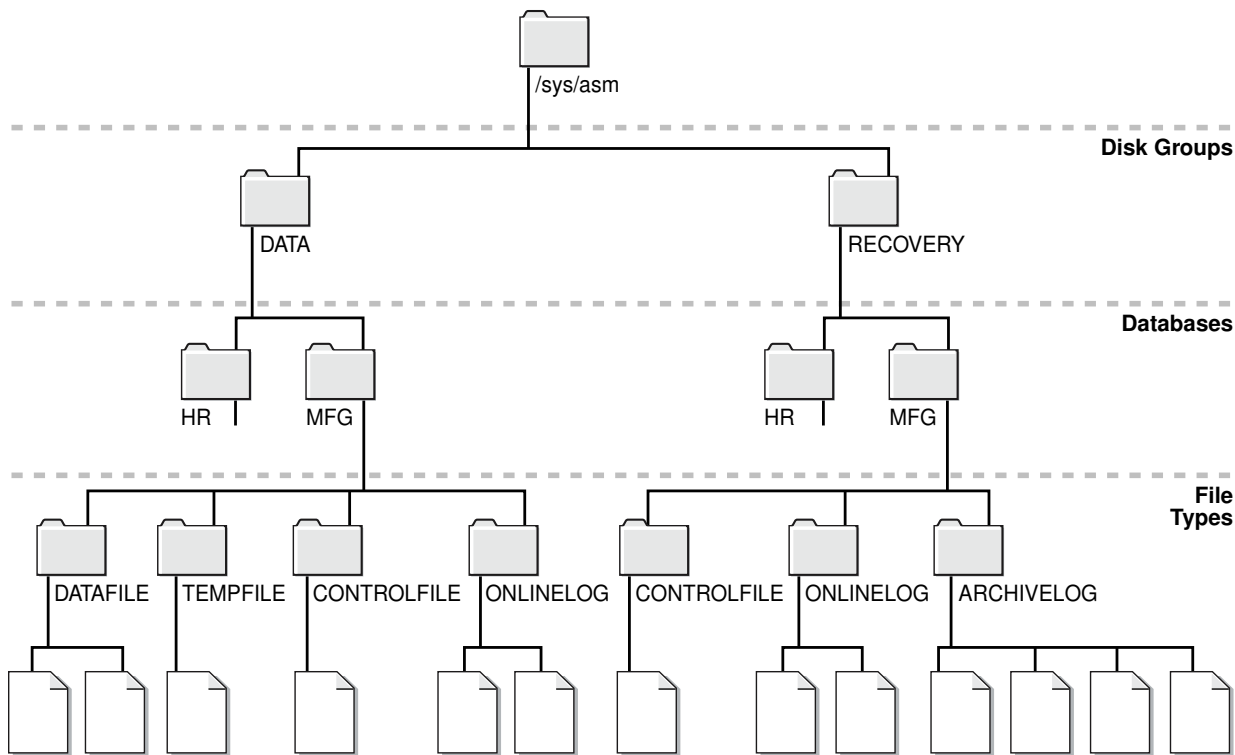
5.6.1 Inside /sys/asm

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 5-1](#) (page 5-18) 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 5-1 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.

5.7 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

5.8 Managing Disk Group 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](#) (page 4-13), this setting 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 5-6](#) (page 5-21) 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` SQL 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.

This section describes how to manage disk group templates under the following topics:

- [Template Attributes](#) (page 5-20)
- [Adding Templates to a Disk Group](#) (page 5-22)
- [Modifying a Disk Group Template](#) (page 5-23)
- [Dropping Templates from a Disk Group](#) (page 5-23)
- [Creating Tablespaces in Oracle ASM: Specifying Attributes with Templates](#) (page 5-23)

See Also:

- *Oracle Database SQL Language Reference* for information about the ALTER DISKGROUP statement
 - [Example 6-10](#) (page 6-6) for an example of the information displayed in the V\$ASM_TEMPLATE view
 - *Oracle Database Reference* for information about the V\$ASM_TEMPLATE view which lists all of the templates known to the Oracle ASM instance
-
-

5.8.1 Template Attributes

[Table 5-4](#) (page 5-20) shows the permitted striping attribute values and allocation unit size chunks. These values correspond to the STRIPE column of V\$ASM_TEMPLATE.

Table 5-4 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 5-5](#) (page 5-20) shows the permitted redundancy values for Oracle ASM templates. These values correspond to the REDUND column of V\$ASM_TEMPLATE.

Table 5-5 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 5-6](#) (page 5-21) shows the initial attribute settings for some default templates. The type of mirroring associated with the **Mirroring** column for normal, high, and external redundancy disk groups is specified in [Table 5-5](#) (page 5-20). For example, the **Mirror** column for high redundancy disk groups displays MIRROR in [Table 5-6](#) (page 5-21). In [Table 5-5](#) (page 5-20), the MIRROR value for high redundancy disk groups corresponds to Three-way mirroring.

For an example of the output of the V\$ASM_TEMPLATE view, refer to [Example 6-10](#) (page 6-6).

Table 5-6 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
ARCHIVELOG	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
ASMPARAMETERFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
AUDIT_SPILLFILES	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
AUTOBACKUP	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
BACKUPSET	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
CHANGETRACKING	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
CONTROLFILE	FINE	HIGH	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
DATAFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
DATAGUARDCONFIG	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
DUMPSET	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
FLASHBACK	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
FLASHFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
KEY_STORE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
OCRFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
ONLINELOG	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
PARAMETERFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
TEMPFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
VOTINGFILE	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD
XTRANSPORT	COARSE	MIRROR	MIRROR	UNPROTECTED	COLD	MIRRORCOLD

[Table 5-7](#) (page 5-22) 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](#) (page 4-26)".

Table 5-7 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.

5.8.2 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 5-6](#) (page 5-21), 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, then 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 5-14](#) (page 5-22) creates a new template named `reliable` for the normal redundancy disk group data.

[Example 5-15](#) (page 5-22) creates a new template named `unreliable` that specifies `UNPROTECTED` (no mirroring) for files.

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.

Example 5-14 Adding a template with high redundancy to a disk group

```
ALTER DISKGROUP data ADD TEMPLATE reliable ATTRIBUTES (HIGH FINE);
```

Example 5-15 Adding a template with external redundancy to a disk group

```
ALTER DISKGROUP data ADD TEMPLATE unreliable ATTRIBUTES (UNPROTECTED);
```

See Also:

Oracle Database SQL Language Reference for more information about the `ALTER DISKGROUP...ADD TEMPLATE` command

5.8.3 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 5-16](#) (page 5-23) changes the striping attribute specification of the `reliable` template for disk group `data`.

Example 5-16 Modifying a disk group template

```
ALTER DISKGROUP data MODIFY TEMPLATE reliable
    ATTRIBUTES (COARSE);
```

5.8.4 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 5-17](#) (page 5-23) drops the previously created template `unreliable` from `data`:

Example 5-17 Dropping a template from a disk group

```
ALTER DISKGROUP data DROP TEMPLATE unreliable;
```

5.8.5 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 5-18](#) (page 5-23) assumes that the template (`mytemplate`) has been defined.

Example 5-18 Using a user-defined template to specify attributes

```
ALTER SYSTEM SET DB_CREATE_FILE_DEST = '+data(mytemplate)';

CREATE TABLESPACE mytblspace;
```

Using Views to Display Oracle ASM Information

Dynamic views display important information about Oracle ASM.

This chapter contains information about using dynamic views to display Oracle Automatic Storage Management (Oracle ASM) information.

- [Views Containing Oracle ASM Disk Group Information](#) (page 6-1)
- [Viewing Oracle ASM File Access Control Information](#) (page 6-5)
- [Viewing Disk Region Information](#) (page 6-6)

See Also:

Oracle Database Reference for information about all of the V\$ASM* dynamic performance views

6.1 Views Containing Oracle ASM Disk Group Information

You can use the views in [Table 6-1](#) (page 6-1) to obtain information about Oracle ASM disk groups.

Note:

The REDUNDANCY_LOWERED column in the V\$ASM_FILE view no longer provides information about files with reduced redundancy. This column is deprecated, and it always displays a value of U.

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_AUDIT_CLEAN_EVENTS	Displays information about the history of audit trail cleanup or purge events.

Table 6-1 (Cont.) Oracle ASM dynamic views for disk group information

View	Description
V\$ASM_AUDIT_CLEANUP_JOBS	Displays information about the configured audit trail purge jobs.
V\$ASM_AUDIT_CONFIG_PARAMS	Displays information about the currently configured audit trail properties.
V\$ASM_AUDIT_LAST_ARCH_TS	Displays information about the last archive timestamps set for audit trail cleanup or purges.
V\$ASM_CLIENT	In an Oracle ASM instance, identifies databases using disk groups managed by the Oracle ASM instance. In an Oracle Database 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 an Oracle Database instance, only the rows for that instance are shown.
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, this view does not perform a discovery when it is queried. To see more complete information on all disks, use V\$ASM_DISKGROUP.
V\$ASM_ESTIMATE	Displays an estimate of the work involved in execution plans for Oracle ASM disk group rebalance and resync operations.
V\$ASM_FILE	Contains one row for every Oracle ASM file in every disk group mounted by the Oracle ASM instance.
V\$ASM_FILEGROUP	Contains one row for every file group in every disk group mounted by the Oracle ASM instance.
V\$ASM_FILEGROUP_FILE	Contains one row for each file associated with a file group in every disk group mounted by the Oracle ASM instance.
V\$ASM_FILEGROUP_PROPERTY	Contains one row for every property associated with the file group or each file type of every file group 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 an Oracle Database instance, contains no rows.

Table 6-1 (Cont.) Oracle ASM dynamic views for disk group information

View	Description
V\$ASM_QUOTAGROUP	Contains one row for every Oracle Automatic Storage Management (Oracle ASM) quota group discovered by the Oracle ASM instance on the node.
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, 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.

Example 6-1 Viewing disk group attributes with V\$ASM_ATTRIBUTE

An example of the use of the V\$ASM_ATTRIBUTE and V\$ASM_DISKGROUP views is shown in this example. 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.

```
SELECT SUBSTR(dg.name,1,12) AS diskgroup, SUBSTR(a.name,1,24) 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
       AND a.name NOT LIKE '%template%';
```

DISKGROUP	NAME	VALUE	READ_ON
DATA	thin_provisioned	FALSE	N
DATA	access_control.umask	066	N
DATA	phys_meta_replicated	true	Y
DATA	disk_repair_time	3.6h	N
DATA	idp.boundary	auto	N
DATA	idp.type	dynamic	N
DATA	content.type	data	N
DATA	content.check	FALSE	N
DATA	au_size	1048576	Y
DATA	sector_size	512	Y
DATA	compatible.asm	12.1.0.0.0	N
DATA	compatible.rdbms	12.1.0.0.0	N
DATA	compatible.advm	12.1.0.0.0	N
DATA	cell.smart_scan_capable	FALSE	N
DATA	access_control.enabled	FALSE	N
DATA	failgroup_repair_time	24.0h	N

Example 6-2 Viewing the compatibility of a disk group with V\$ASM_DISKGROUP

You can view the compatibility for a disk group with the V\$ASM_DISKGROUP view, as shown in this example.

```
SELECT name AS diskgroup, substr(compatibility,1,12) AS asm_compat,
       substr(database_compatibility,1,12) AS db_compat FROM V$ASM_DISKGROUP;
```

DISKGROUP	ASM_COMPAT	DB_COMPAT
DATA	12.1.0.0.0	12.1.0.0.0
FRA	12.1.0.0.0	12.1.0.0.0

Example 6-3 Viewing disks in disk groups with V\$ASM_DISK

An example of the use of the V\$ASM_DISK and V\$ASM_DISKGROUP views is shown in this example. This example displays the disks associated with a disk group, plus the mount status, state, and failure group of the disks.

```
SELECT SUBSTR(dg.name,1,16) AS diskgroup, SUBSTR(d.name,1,16) AS asmdisk,
       d.mount_status, d.state, SUBSTR(d.failgroup,1,16) AS failgroup
FROM V$ASM_DISKGROUP dg, V$ASM_DISK d WHERE dg.group_number = d.group_number;
```

DISKGROUP	ASMDISK	MOUNT_S	STATE	FAILGROUP
DATA	DATA_0000	CACHED	NORMAL	DATA_0000
DATA	DATA_0010	CACHED	NORMAL	DATA_0010
DATA	DATA_0001	CACHED	NORMAL	DATA_0001
DATA	DATA_0003	CACHED	NORMAL	DATA_0003
DATA	DATA_0009	CACHED	NORMAL	DATA_0009
DATA	DATA_0007	CACHED	NORMAL	DATA_0007
DATA	DATA_0004	CACHED	NORMAL	DATA_0004
DATA	DATA_0008	CACHED	NORMAL	DATA_0008
DATA	DATA_0006	CACHED	NORMAL	DATA_0006
DATA	DATA_0011	CACHED	NORMAL	DATA_0011
DATA	DATA_0005	CACHED	NORMAL	DATA_0005
DATA	DATA_0002	CACHED	NORMAL	DATA_0002
FRA	FRA_0011	CACHED	NORMAL	FRA_0011
FRA	FRA_0002	CACHED	NORMAL	FRA_0002
FRA	FRA_0001	CACHED	NORMAL	FRA_0001
FRA	FRA_0003	CACHED	NORMAL	FRA_0003
FRA	FRA_0004	CACHED	NORMAL	FRA_0004
FRA	FRA_0000	CACHED	NORMAL	FRA_0000
...				

Example 6-4 Viewing disk group clients with V\$ASM_CLIENT

An example of the use of the V\$ASM_CLIENT and V\$ASM_DISKGROUP views on an Oracle ASM instance is shown in this example, which displays disk groups with information about the connected database client instances.

```
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	12.1.0.1.0	12.1.0.1.0
DATA	orcl	orcl	12.1.0.1.0	12.1.0.1.0
DATA	+ASM	asmvol	12.1.0.1.0	12.1.0.1.0
FRA	orcl	orcl	12.1.0.1.0	12.1.0.1.0
...				

6.2 Viewing Oracle ASM File Access Control Information

You can view information about Oracle ASM File Access Control in the columns of the V\$ASM_USER, V\$ASM_USERGROUP, V\$ASM_USERGROUP_MEMBER, and V\$ASM_FILE views.

For more information about Oracle ASM File Access Control, see "[Managing Oracle ASM File Access Control for Disk Groups](#) (page 4-61)".

Example 6-5 Viewing Oracle ASM File Access Control information with V\$ASM_USER

This example shows information about Oracle ASM File Access Control users displayed in the V\$ASM_USER view.

```
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 Viewing File Access Control information with V\$ASM_USERGROUP

This example shows information about Oracle ASM File Access Control user groups displayed in the V\$ASM_USERGROUP view.

```
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 Viewing File Access Control information with V\$ASM_USERGROUP_MEMBER

This example shows information about Oracle ASM File Access Control user groups and members displayed in the V\$ASM_USERGROUP_MEMBER view.

```
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 Viewing Oracle ASM File Access Control information with V\$ASM_FILE

This example shows information about Oracle ASM File Access Control file permissions displayed in the V\$ASM_FILE view.

```
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.name = 'FRA' AND dg.group_number = u.group_number AND
      u.group_number = ug.group_number AND ug.group_number = f.group_number AND
      f.group_number = a.group_number 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
...						

6.3 Viewing Disk Region Information

Note:

The Oracle ASM Intelligent Data Placement (IDP) feature has been deprecated in Oracle ASM 12c Release 2 (12.2) and may be desupported in a future release.

Information about Intelligent Data Placement is displayed in the columns of the V\$ASM_DISK, V\$ASM_DISK_IOSTAT, V\$ASM_FILE, and V\$ASM_TEMPLATE views.

For information about setting Intelligent Data Placement, see "[Intelligent Data Placement](#) (page 4-26)".

Example 6-9 Viewing Intelligent Data Placement information with V\$ASM_FILE

This example shows queries for Intelligent Data Placement information in the PRIMARY_REGION and MIRROR_REGION columns of the V\$ASM_FILE view.

```
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 Viewing Intelligent Data Placement information with V\$ASM_TEMPLATE

This example displays Intelligent Data Placement information in the PRIMARY_REGION and MIRROR_REGION columns of the V\$ASM_TEMPLATE view.

```
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	AUDIT_SPILLFILES	COARSE	MIRROR	COLD	COLD

DATA	AUTOBACKUP	COARSE MIRROR COLD COLD
DATA	AUTOLOGIN_KEY_STORE	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	FLASHFILE	COARSE MIRROR COLD COLD
DATA	INCR XTRANSPORT BACKUPSET	COARSE MIRROR COLD COLD
DATA	KEY_STORE	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	VOTINGFILE	COARSE MIRROR COLD COLD
DATA	XTRANSPORT	COARSE MIRROR COLD COLD
DATA	XTRANSPORT BACKUPSET	COARSE MIRROR COLD COLD

22 rows selected.

Administering Oracle ASM with Oracle Enterprise Manager

Oracle Enterprise Manager Cloud Control provides tools for administering Oracle ASM.

This chapter describes how to administer Oracle Automatic Storage Management (Oracle ASM) with Oracle Enterprise Manager Cloud Control.

This chapter contains the following topics:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)
- [Configuring Oracle ASM Initialization Parameters with Oracle Enterprise Manager](#) (page 7-2)
- [Managing Oracle ASM Users with Oracle Enterprise Manager](#) (page 7-3)
- [Managing Disk Groups with Oracle Enterprise Manager](#) (page 7-4)
- [Managing Oracle ASM File Access Control with Oracle Enterprise Manager](#) (page 7-18)
- [Managing Directories, Files, and Aliases with Oracle Enterprise Manager](#) (page 7-19)
- [Managing Disk Group Templates with Oracle Enterprise Manager](#) (page 7-19)
- [Managing Oracle ASM Filter Driver With Oracle Enterprise Manager](#) (page 7-20)
- [Monitoring Oracle ASM Performance With Oracle Enterprise Manager](#) (page 7-21)
- [Backing Up Oracle ASM Files with Oracle Enterprise Manager](#) (page 7-22)
- [Performing Bad Block Recovery with Oracle Enterprise Manager](#) (page 7-23)
- [Migrating to Oracle ASM with Oracle Enterprise Manager](#) (page 7-23)
- [Monitoring Alerts and Incidents With Oracle ASM Support Workbench](#) (page 7-25)
- [Monitoring Oracle Flex ASM](#) (page 7-25)

See Also:

- *Oracle Enterprise Manager Cloud Administration Guide*
 - *Oracle Enterprise Manager Administration*
-
-

7.1 Accessing the Oracle Automatic Storage Management Home Page

All Oracle ASM administration tasks begin with the Oracle Automatic Storage Management home page in Oracle Enterprise Manager Cloud Control. 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 System (Oracle ACFS) 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.

To access the Oracle Automatic Storage Management home page on a single-instance system:

1. Log in to Oracle Enterprise Manager Cloud Control.
2. Under the **Targets** drop down menu at the top of the page, select **All Targets**.
3. In the lists of targets on the **All Targets** page, click the target for the Oracle ASM instance.
4. 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 `SYSASM`. The Oracle Automatic Storage Management home page displays.

For more information about authentication, refer to "[Authentication for Accessing Oracle ASM Instances](#) (page 3-27)".

7.2 Configuring Oracle ASM Initialization Parameters with Oracle Enterprise Manager

You can configure Oracle ASM with the Oracle Enterprise Manager Configuration Parameters page.

To configure the Oracle ASM instance:

1. Access the Oracle Automatic Storage Management home page.

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 as follows:
 - Disk Discovery Path (ASM_DISKSTRING initialization parameter)
Enter a string value.
 - Auto Mount Disk Groups (ASM_DISKGROUPS initialization parameter)
Enter a string value.
 - Rebalance Power (ASM_POWER_LIMIT initialization parameter)
Enter an integer value.
 - 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.
4. Click **Apply** to save your changes or click **Revert** to discard any changes made on the page.

See Also:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for instructions on accessing the Oracle Automatic Storage Management home page
 - [ASM_DISKSTRING](#) (page 3-10) for information about the ASM_DISKSTRING initialization parameter
 - [ASM_DISKGROUPS](#) (page 3-9) for information about the ASM_DISKGROUPS initialization parameter
 - [ASM_POWER_LIMIT](#) (page 3-11) for information about the ASM_POWER_LIMIT initialization parameter
 - [ASM_PREFERRED_READ_FAILURE_GROUPS](#) (page 3-12) for information about the ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter
-

7.3 Managing Oracle ASM Users with Oracle Enterprise Manager

You can manage Oracle ASM users with the Oracle Enterprise Manager Users page . Access this page by clicking the Users tab that is near the top center of the Oracle Automatic Storage Management home page.

To manage Oracle ASM users:

1. Access the Oracle Automatic Storage Management home page.
2. Click the **Users** link at the top to view the Users property page.

ASMSNMP is a less privileged user that is primarily used by Oracle Enterprise Manager to monitor Oracle ASM instances. This account is granted the SYSDBA privilege. ASMSNMP is created by Oracle ASM Configuration Assistant (ASMCA) when an Oracle ASM instance is created.

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.

3. To edit a user's properties, select the box in the Select column of the user to edit, then 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.
4. To delete an Oracle ASM user, select the box in the Select column for the user to delete, then click **Delete**. On the confirmation dialog, click **Yes** to delete the user or **No** to stop the user deletion procedure.

See Also:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for instructions on accessing the Oracle Automatic Storage Management home page
 - [Creating Oracle ASM Instances with ASMCA](#) (page 9-2) for information about creating an Oracle ASM instance with ASMCA
 - [Authentication for Accessing Oracle ASM Instances](#) (page 3-27) for information about authentication for Oracle ASM instances
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for information about installing Oracle ASM
 - *Oracle Database Installation Guide* for information about reviewing database accounts and passwords
-
-

7.4 Managing Disk Groups with Oracle Enterprise Manager

This section describes how to manage disk groups under the following topics:

- [Creating Disk Groups](#) (page 7-5)
- [Adding Disks to Disk Groups](#) (page 7-7)
- [Dropping Disks from Disk Groups](#) (page 7-9)

- [Bringing Disks Online and Offline](#) (page 7-8)
- [Dropping Disk Groups](#) (page 7-11)
- [Monitoring Disk Group Usage](#) (page 7-12)
- [Mounting and Dismounting Disk Groups](#) (page 7-12)
- [Administering Advanced Disk Group Properties](#) (page 7-13)
- [Specifying a Power Limit For Disk Online](#) (page 7-14)
- [Replacing An Existing Disk With a New Disk](#) (page 7-15)
- [Setting and Modifying Content Type for the Disk Group](#) (page 7-15)
- [Setting and Modifying the Failure Group Repair Time](#) (page 7-16)
- [Running Rebalance on Multiple Disk Groups](#) (page 7-16)
- [Viewing Rebalance Process Related Data](#) (page 7-17)
- [Setting and Modifying the Thin Provisioned Disk Group Attribute](#) (page 7-17)
- [Running the On Demand Scrubbing Process](#) (page 7-17)
- [Determining Whether Scrubbing Process Is Running](#) (page 7-18)

For more information about managing Oracle ASM disk groups, refer to [Administering Oracle ASM Disk Groups](#) (page 4-1)

7.4.1 Creating Disk Groups

You can create disk groups with Oracle Enterprise Manager.

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.
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. 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.

5. 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), 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 to be included in the new disk group.
 - Optionally, you can enter a disk name, the size of the disk for Oracle ASM to use, and the failure group name.
 - To force the inclusion of a disk in the disk group, select the **Force Usage** box for that disk.

Caution:

Selecting 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. The existing data is then lost. You must be certain that you are selecting a disk that can legitimately be added to the disk group.

8. Specify a **Disk Compatibility** value for the disk group to enable your system to use the latest Oracle ASM features.

Note:

You cannot reverse the setting of a disk group compatibility value.

9. 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:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for instructions on accessing the Oracle Automatic Storage Management home page
 - [Creating Disk Groups](#) (page 4-9) for information about using the FORCE option when creating disk groups
 - [Configuring Disk Group Compatibility Attributes](#) (page 7-13) for information about configuring disk group compatibility attributes
 - [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility attributes
 - [Creating Disk Groups](#) (page 4-9) for more information about creating disk groups
-
-

7.4.2 Adding Disks to Disk Groups

You can add disks to disk groups with Oracle Enterprise Manager.

You can add disks to a disk group to increase the 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 with the setting of the Rebalance Power. 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. Manual or automatic rebalance can then occur later. You might want to postpone rebalancing until there are fewer demands on the database, or when you have completed all of your disk group changes.

To add one or more disks to a disk group:

1. Access the Oracle Automatic Storage Management home page.
2. 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 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.
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:
 - Check the box to the left of each disk to be added to the disk group.
 - 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. The existing data is then lost. You must be certain that you are selecting a disk that can legitimately be added to the disk group.

- 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:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for instructions on accessing the Oracle Automatic Storage Management home page
 - [Manually Rebalancing Disk Groups](#) (page 4-28) and [Tuning Rebalance Operations](#) (page 4-31) for information about controlling rebalance operations
 - [ASM_POWER_LIMIT](#) (page 3-11) for information about rebalance power settings and the ASM_POWER_LIMIT initialization parameter
 - [Adding Disks to a Disk Group](#) (page 4-22) for an example of using the FORCE option
 - [Altering Disk Groups](#) (page 4-19) for more information about modifying disk group properties
-
-

7.4.3 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 "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
2. Click the **Disk Groups** link at the top of the page to view the Disk Groups page.
Click the name of a disk in the Name column and Oracle Enterprise Manager displays the General page for the disk group to modify.
3. To bring a disk online, select the disk 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 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.
4. To take a disk offline, select the disk 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. 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. For information about rebalance power settings and the `ASM_POWER_LIMIT` initialization parameter, refer to "[ASM_POWER_LIMIT](#) (page 3-11)".
 - 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 time interval 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.

7.4.4 Dropping Disks from Disk Groups

You can drop disks from disk groups with Oracle Enterprise Manager.

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
- To upgrade a disk
- To reallocate the disk to a different disk group or reallocate the disk 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.
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 that contains the disks to be dropped.

The Disk Group page appears.

5. Under the Member Disks heading, in the Select column, select the boxes for the disks that you plan to drop, and then click **Remove**.

A confirmation page appears.

6. To change the rebalance power, or to 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.

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 Also:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for instructions about accessing the Oracle Automatic Storage Management home page
 - [Dropping Disks from Disk Groups](#) (page 4-25) for information about using the `FORCE` option when dropping disk groups
 - [Dropping Disks from Disk Groups](#) (page 4-25) for more information about dropping disks
-
-

7.4.5 Dropping Disk Groups

When you drop a disk group, you delete the disk group optionally with 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 "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
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 plan to drop.
5. Click **Delete** and the confirmation page appears.
6. 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**.
For more information about dropping disks, see "[Dropping Disk Groups](#) (page 4-69)".

7.4.6 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 "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
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 displays all of the disk groups and their space usage information.

The Usable Free column displays the space in megabytes that is 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.

For more information, see "[Managing Capacity in Disk Groups](#) (page 4-35)" and "[Views Containing Oracle ASM Disk Group Information](#) (page 6-1)"

7.4.7 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 plan to force mount, then click **Mount**.
5. On the Disk Mount Confirmation dialog, click the **Show Advanced Options** link.
6. On the Mount Force Confirmation dialog, 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.

7.4.8 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](#) (page 7-13)
- [Configuring Disk Repair Time](#) (page 7-14)
- [Configuring Smart Scan Compatibility](#) (page 7-14)
- [Configuring File Access Control](#) (page 7-14)

7.4.8.1 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 the setting of a disk group 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 ASM Dynamic Volume Manager (Oracle ADVM) functionality. The ASM compatibility of the disk group must be set to 11.2 or greater to enable Oracle ADVM. Also, the Oracle ADVM volume drivers must be loaded in the supported environment.

For information about disk group compatibility attributes, see "[Disk Group Compatibility](#) (page 4-50)".

7.4.8.2 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](#) (page 4-43)".

7.4.8.3 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

7.4.8.4 Configuring File Access Control

Specify to disable or enable Oracle ASM File Access Control. If Oracle ASM File Access Control is enabled, then 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](#) (page 4-61)".

7.4.9 Specifying a Power Limit For Disk Online

To specify a power limit for disk online, follow these steps:

1. Access the Oracle Automatic Storage Management home page.
2. Select the **Disk Groups** tab. The Disk Groups page is displayed.
3. Select a disk group. The Disk Group page appears and shows all Member Disks that belong to the disk group in the Member Disks table.
4. Select a disk, then click **Online**. The online confirmation page is displayed with the advance option to specify a power limit.

5. Specify the **Power Limit**, then click **Submit**. The disk is set to online with the specified power limit.

See Also:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for instructions about accessing the Oracle ASM home page
 - [Manually Rebalancing Disk Groups](#) (page 4-28) for information about the power option in rebalancing a disk group
 - [ASM_POWER_LIMIT](#) (page 3-11) for information about the ASM_POWER_LIMIT initialization parameter
-
-

7.4.10 Replacing An Existing Disk With a New Disk

To replace an existing disk with a new disk, follow these steps:

1. Access the Oracle Automatic Storage Management home page.

See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.

2. Click the Disk Group tab. The Disk Group page displays.

3. Select a Disk Group. The Disk Group page appears which shows all Member Disks that belong to the disk group.

4. Select an offline disk to replace. The Replace Disks page is displayed.

Select a new disk of size equal to or greater than the original disk being replaced to ensure the same level of data protection.

5. Select a new disk and then specify the options, such as Rebalance Power, and then click **OK**. The disk is then replaced. Rebalance operations redistribute data evenly across all drives. Oracle 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 1024.

For additional information about replacing a disk, refer to "[Replacing Disks in Disk Groups](#) (page 4-24)".

7.4.11 Setting and Modifying Content Type for the Disk Group

To set and modify the content type for a disk group, follow these steps:

1. Access the Oracle Automatic Storage Management home page.

See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.

2. Click the Disk Groups tab. The Disk Group page displays.

3. Click **Create**. In the Disk Resync Attributes section of the Create Disk Group page, you can select the value **Recovery**, **Data** or **System** from the **Content Type** drop-

down list. The attribute can be set only if the ASM compatibility is greater than or equal to 12.1.0.0.

4. To edit the content type for an existing disk group, navigate to the Disk Group home page by clicking the disk group in the Disk Groups table. The Disk Group home page displays.
5. Click **Edit** in the Advanced Attribute Section. The Edit Advanced Attributes for Disk Group page appears. In the Disk Resync Attributes section, select from the values **Recovery**, **Data**, or **System** from the **Content Type** drop-down list.

For information about the `CONTENT . TYPE` disk group attribute, refer to "[CONTENT . TYPE](#) (page 4-5)".

7.4.12 Setting and Modifying the Failure Group Repair Time

To set and modify the Failure Group Repair Time, follow these steps:

1. Access the Oracle Automatic Storage Management home page.

See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
2. Click the Disk Groups tab. The Disk Groups page displays.
3. Click **Create**. Specify the repair time for the **Failure Group** in the Candidate Member Disks table.
4. To edit an existing Failure Group repair time, from the Disk Group page choose the **Disk Group** for the repair time you want to modify. The Disk Group page appears.
5. In the Advanced Attributes section, click **Edit**. The Edit Advanced Attributes for Disk Group page displays.
6. In the Disk Resync Attributes section, modify the **Failure Group Repair Time**, then click **OK**. The failure Group Repair Time applies to all the failure groups in the disk group except those containing only one disk. This value overrides the disk repair time and is overridden by the actual value provided when a disk is taken offline.

For additional information about setting and modifying the failure group repair time, refer to "[Oracle ASM Fast Mirror Resync](#) (page 4-43)".

7.4.13 Running Rebalance on Multiple Disk Groups

Enterprise Manager enables you to rebalance on multiple disk groups and distribute the data copy across other member instances in the cluster. To run rebalance on multiple disk groups, follow these steps:

1. Access the Oracle Automatic Storage Management home page.

See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
2. Click the Disk Group tab. The Disk Group page displays.
3. Select multiple disk groups then click **Rebalance**. The Rebalance Disk Groups Confirmation page appears where you can set the Rebalance Power by displaying Advanced Options.

4. Set the Rebalance Power value, then click **OK**. The command to rebalance is submitted.

For additional information about rebalancing disk groups, refer to "[Manually Rebalancing Disk Groups](#) (page 4-28)".

7.4.14 Viewing Rebalance Process Related Data

To view rebalance process-related data, including the number of participating instances on the Pending Operations page, follow these steps:

1. Access the Oracle Automatic Storage Management home page.
See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
2. Click the Disk Group tab. The Disk Groups page displays.
3. Click a Disk Group link. The Disk Group home page General section displays the number of **Pending Operations** as a link to the Pending Operations page.
4. Click the link. Information on the rebalance process, along with details on the number of participating instances, are displayed.

For additional information about rebalancing disk groups, refer to "[Tuning Rebalance Operations](#) (page 4-31)".

7.4.15 Setting and Modifying the Thin Provisioned Disk Group Attribute

To set and modify the thin-provisioned disk group attribute, follow these steps:

1. Access the Oracle Automatic Storage Management home page.
See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
2. Click the Disk Groups tab. The Disk Groups page is displayed.
3. Click the Disk Group that contains the Thin Provisioned attribute you want to change. The Disk Group page displays.
4. In the Advance Attributes section, click **Edit**. The Edit Advanced Attributes for Disk Group page displays.
5. In the Thin Provisioning section, select **True** or **False** from the Thin Provisioned field drop-down list. If thin provisioning is enabled, the disk space not in use can be returned to the array. This attribute can be set only if the ASM compatibility is greater than or equal to 12.1.0.0.

For information about the `THIN_PROVISIONED` disk group attribute, refer to "[THIN_PROVISIONED](#) (page 4-9)".

7.4.16 Running the On Demand Scrubbing Process

To run the On Demand scrubbing process, follow these steps:

1. Access the Oracle Automatic Storage Management home page.

See "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)" for instructions.

2. Click the Disk Group tab. The Disk Group page is displayed.
3. Click **Scrub**. The Scrub Disk Group page is displayed where you can use advanced options such as repair, power level, wait and force. You can also use the scrub action on disks and files.
4. Provide the advance option and click **OK**. A background process is initiated that scrubs the disk group.

For additional information about scrubbing disk groups, refer to "[Scrubbing Disk Groups \(page 4-32\)](#)".

7.4.17 Determining Whether Scrubbing Process Is Running

To determine whether the scrubbing process is running, follow these steps:

1. Access the Oracle Automatic Storage Management home page.

See "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)" for instructions.
2. Click the Disk Group tab. The Disk Group page displays.
3. Click a Disk Group link. The Disk Group General Section displays a Pending Operation link if a scrubbing process is ongoing.
4. Click the link. Details on the scrubbing process are displayed.

7.5 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.

To manage Oracle ASM File Access Control:

1. Access the Oracle Automatic Storage Management home page.

See "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)" 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, remove, or replace users from a user group and create or delete user groups. You can also modify the ownership or permission of files.

To access files in a disk group, operating system users must be granted access privileges to the disk group. After granted access privileges, users can be added as members of groups. The operating system user who owns the oracle binary for the database instance is automatically added and cannot be removed.

Oracle ASM File Access Control operations on this page are available when the files are open.

For more information about Oracle ASM File Access Control, see "[Managing Oracle ASM File Access Control for Disk Groups](#) (page 4-61)".

7.6 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.

To manage Oracle ASM files:

1. Access the Oracle Automatic Storage Management home page.
See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" 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 icons to left on the folders to expand or contract the folders. When you expand the folders, you can drill down to individual files.
6. 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.

For more information about managing Oracle ASM files, directories, and aliases, see [Administering Oracle ASM Files, Directories, and Templates](#) (page 5-1).

7.7 Managing Disk Group Templates with Oracle Enterprise Manager

You can manage disk group templates with the Templates tab of the Disk Group page.

To manage Oracle ASM templates:

1. Access the Oracle Automatic Storage Management home page.
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).

see Also:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for instructions about accessing the Oracle ASM home page
 - [Managing Disk Group Templates](#) (page 5-19) for more information about disk group templates
 - [Intelligent Data Placement](#) (page 4-26) for information about hot and cold disk regions
-

7.8 Managing Oracle ASM Filter Driver With Oracle Enterprise Manager

You can manage Oracle ASM Filter Driver with Oracle Enterprise Manager.

This section contains the following sections:

- [Discovering ASM Filter Driver Disks](#) (page 7-20)
- [Provisioning ASM Filter Driver Disks](#) (page 7-20)
- [Releasing ASM Filter Driver Disks](#) (page 7-21)

For more information about managing Oracle ASM Filter Driver, refer to "[Administering Oracle ASM Filter Driver](#) (page 20-1)".

7.8.1 Discovering ASM Filter Driver Disks

To discover disks for Oracle ASM Filter Driver follow these steps:

1. From **Targets** menu, select **All Targets**.
2. On the All Targets page, select any target of type **Automatic Storage Management**.
3. On the target home page, from **Automatic Storage Management** menu, select **Configuration Parameters**.
4. On the Configuration page, click **Manage Disk Provisioning**. The ASM Filter Driver Page is displayed.
5. On the Disks available for Provisioning table, click **Discover** to discover additional Oracle ASM disks.
6. In the Discover Disks dialog box, enter additional discovery paths from where you would like to provision disks. If you are providing more than one path, then ensure that the values are separated by a comma.

For example, ' /dev/ ', ' /tmp '.

7. Click **Ok**.

In the Disks available for Provisioning table, all the disks available in the discovery paths are listed.

7.8.2 Provisioning ASM Filter Driver Disks

To provision a set of disks for Oracle ASM Filter Driver, follow these steps:

1. From **Targets** menu, select **All Targets**.
2. On the All Targets page, select any target of type **Automatic Storage Management**.
3. On the target home page, from **Automatic Storage Management** menu, select **Configuration Parameters**.
4. On the Configuration page, click **Manage Disk Provisioning**. The ASM Filter Driver Page is displayed.
5. From the Disks available for Provisioning table, select one or more disks, then click **Provision**.
6. The Provision Disks dialog box lists all the disks that were selected for provisioning. You can update the Label info if you want, then click **OK**.

After provisioning, the Oracle ASM disks appear in the Provisioned Disks table. These disks display with their new label. For example, if you select `/dev/loop3` to be provisioned with a label `LOOP3`, then after provisioning the disk displays as `AFD:LOOP3` in the Provisioned Disks table.

After the disks are provisioned, they become eligible to be managed by the Oracle ASM Filter Driver.

7.8.3 Releasing ASM Filter Driver Disks

To release the disks currently managed by Oracle ASM Filter Driver, follow these steps:

1. From **Targets** menu, select **All Targets**.
2. On the All Targets page, select any target of type **Automatic Storage Management**.
3. On the target home page, from **Automatic Storage Management** menu, select **Configuration Parameters**.
4. On the Configuration page, click **Manage Disk Provisioning**. The ASM Filter Driver Page is displayed.
5. From the Provisioned Disks table, select one or more disks, then click **Release**.
6. In the Release Provisioned Disks dialog box, click **Yes** to release the disks

After the disks are released, they are removed from the Provisioned Disks table. They instead appear in the Disks available for provisioning table as a candidate eligible to be provisioned.

7.9 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](#) (page 7-22)
- [Checking Disk Group Integrity](#) (page 7-22)

7.9.1 Monitoring General Oracle ASM Performance

You can monitor Oracle ASM performance with the Oracle Enterprise Manager Performance page.

To monitor Oracle ASM performance:

1. Access the Oracle Automatic Storage Management home page.
See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
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

For more information, refer to "[Performance and Scalability Considerations for Disk Groups](#) (page 4-48)".

7.9.2 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** page, select a disk group in the list and click **Check**.
A confirmation dialog for the check disk group operation displays.
2. Select **Check Without Repair** 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** 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. To schedule a job, click **Schedule Job**. 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. Click **OK** to schedule the job or **Cancel** to terminate the job creation process.

7.10 Backing Up Oracle ASM Files with Oracle Enterprise Manager

Oracle recommends that you use RMAN to back up files managed by Oracle ASM.

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.

See Also:

- [Performing Oracle ASM Data Migration with RMAN](#) (page 8-1) for information about migrating data with RMAN
 - *Oracle Database 2 Day DBA* for information about the procedures for using RMAN
 - *Oracle Database Backup and Recovery User's Guide* for additional information about RMAN
-

7.11 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. Access the Oracle Automatic Storage Management home page.

See "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)" for instructions.
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 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 in which to recover bad blocks, then 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.

7.12 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 **Administration** pull down list and select **Migrate to Oracle ASM**.

2. You may need to set credentials for the database, Oracle ASM, and host.

If prompted for database login credentials, then enter the user `SYS`, provide the `SYS` password that was set for the Oracle Database instance during installation, and connect as `SYSDBA`.

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 `SYSASM`. For more information about authentication, see "[Authentication for Accessing Oracle ASM Instances](#) (page 3-27)".

3. 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 plan 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 option. Otherwise, using offline migration is the default. If you perform an online migration, then you can suspend the migration by selecting the check box 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.

4. 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.
5. 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.
6. When you continue, Oracle Enterprise Manager displays the Schedule page on which you can specify a job name and a date and time for the migration job to begin. Click **Next** when you have configured your desired settings.
7. 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.

Note:

If your migration job fails, then do not start your database until you complete a recovery operation. Otherwise, you may corrupt your database.

8. 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 may need to wait for the links to display after a subsequent refresh.

7.13 Monitoring Alerts and Incidents With Oracle ASM Support Workbench

Oracle Enterprise Manager provides Support Workbench to monitor Oracle ASM alerts and incidents.

To access Support Workbench, select **Diagnostics > Support Workbench** under the **Automatic Storage Management** target menu at the top of the page.

On the Oracle ASM Support Workbench page **Problems** tab you can view information about current and past problems. You also have options to view additional details and create a package that you can send to Oracle Support Services.

Support Workbench guides you through the packaging process as follows:

1. To create a package, select an incident and click **Package** on the Support Workbench Problems page. The Support Workbench **Select packaging mode** page of the wizard displays.
2. On the **Select packaging mode** page, select either the **Quick** or **Custom** packaging mode.
3. Depending on the packaging mode chosen, you are presented with various entry fields and options as you click **Next** through the wizard.

The wizard guides you through the packaging process using a series of steps that include:

- Viewing and editing the contents of the package
- Viewing and editing the manifest of the package
- Adding additional diagnostic details
- Scheduling the generation and uploading of the package

On the Support Workbench **Checker Findings** tab, you can view and edit information about packages created for problems on the Support Workbench **Packages** tab. You can also check findings on the **Checker Findings** tab.

See Also:

Oracle Database Administrator's Guide for information about viewing problems with Oracle Enterprise Manager Support Workbench

7.14 Monitoring Oracle Flex ASM

Oracle Enterprise Manager provides functionality to monitor your Oracle Flex ASM configuration.

- On the Oracle Flex ASM home page, you can view and administer components of the Oracle Flex ASM configuration.
- On the Oracle Flex ASM Members page, you can view all the members of the Oracle Flex ASM configuration.
- On the Oracle Flex ASM Topology page, you can view the topology of the Oracle Flex ASM configuration.

For information about Oracle Flex ASM, refer to "[Managing Oracle Flex ASM](#) (page 19-1)".

Performing Oracle ASM Data Migration with RMAN

Recovery Manager (RMAN) provides tools to migrate data to and from Oracle ASM. 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](#) (page 8-1)
- [Preparing to Migrate the Database to Oracle ASM Using RMAN](#) (page 8-3)
- [Migrating the Database to Oracle ASM Using RMAN](#) (page 8-6)
- [Migrating a Database from Oracle ASM to Alternative Storage](#) (page 8-9)
- [Moving Data Files Between Oracle ASM Disk Groups Using RMAN](#) (page 8-10)

The procedures in this chapter cover an Oracle ASM standalone environment on a Linux platform. File locations and procedures may be different for different Oracle configurations and on other operating system platforms.

See Also:

- *Oracle Database Backup and Recovery User's Guide* for complete information about using RMAN
 - Operating system-specific documentation for migrating data in and out of Oracle ASM on specific platforms
 - *Oracle Clusterware Administration and Deployment Guide* for information about managing and migrating Oracle Cluster Registry and voting files
 - *Oracle Real Applications Cluster Administration and Deployment Guide* for migrating data in and out of Oracle ASM in an Oracle RAC configuration
 - Articles at [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com) for information about gathering and backing up Oracle ASM and Oracle ACFS metadata in a formatted manner, such as HTML format
-

8.1 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](#) (page 8-2)
- [Basic Concepts of Oracle ASM Data Migration](#) (page 8-2)
- [Basic Steps of Data Migration to Oracle ASM Using RMAN](#) (page 8-3)

8.1.1 Purpose of Oracle ASM Data Migration

Alternatives to Oracle ASM storage include file systems, raw disks, and SAN configurations. Oracle ASM includes numerous benefits over these storage alternatives, including performance optimization, redundancy protection, and load balancing. You do not need 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.

8.1.2 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. It is not necessary 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 internal database identifier (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.

8.1.3 Basic 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 System Change Number (SCN) with 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](#) (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](#) (page 8-6)".

To migrate files from Oracle ASM to alternative storage, see "[Migrating a Database from Oracle ASM to Alternative Storage](#) (page 8-9)".

8.2 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`.

To prepare the database for Oracle ASM migration:

1. If the `COMPATIBLE` initialization parameter 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 SYSBACKUP privileges (rather than SYSDBA privileges) to enforce the separation of duty security model, and run 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. Connect with SYSBACKUP privileges to enforce the separation of duty security model.

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 a 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 fast 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;
```

8.3 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](#) (page 8-3). If you are not migrating the recovery area to Oracle ASM, then you must modify some 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](#) (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. Run the `CREATE SPFILE` command in SQL*Plus as follows, where *sid* is the SID of the database:

```
SQL> CREATE SPFILE='+DATA/spfilesid.ora' FROM PFILE='?/dbs/initsid.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.

Run 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 "Back up the data files to the Oracle ASM disk group" in [Preparing to Migrate the Database to Oracle ASM Using RMAN](#) (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 enable Flashback Database unless you had disabled it previously.

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 "Set Oracle Managed Files initialization parameters to Oracle ASM locations" in [Migrating the Database to Oracle ASM Using RMAN](#) (page 8-6) are set.

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. Run the following command at the RMAN prompt:

```
RMAN> DELETE ARCHIVELOG ALL;
```

- b. Back up archived redo log files, backup sets, and data file copies to Oracle ASM. For example, run 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;
    }

```

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;
/

```

8.4 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](#) (page 8-3)" and "[Migrating the Database to Oracle ASM Using RMAN](#) (page 8-6)" 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 file name used in a SQL statement is an Oracle ASM location, then change it to a file name in the alternative storage location.

8.5 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. You can use RMAN 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.

You can also use the ALTER DATABASE MOVE DATAFILE SQL statement to move data files. For information about moving data files online with ALTER DATABASE MOVE DATAFILE, refer to "[Moving Data Files Between Disk Groups With ALTER DATABASE](#) (page 5-11)".

For this scenario using RMAN, assume that you are using disk groups DATA and USERDATA and you want to move the data file 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 using the RMAN BACKUP AS COPY procedure with the SET NEWNAME and SWITCH commands, 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.

Run 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.

Run 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.

Run the following SQL command in the RMAN client. 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.

Run the SQL command in the RMAN client. 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. Run 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
```

Managing Oracle ASM With ASMCA

Oracle ASM Configuration Assistant (ASMCA) provides utilities for managing Oracle ASM.

Oracle ASM Configuration Assistant (ASMCA) supports installing and configuring Oracle ASM instances, Oracle ASM disk groups, Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes, and Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems. In addition, you can use the ASMCA command-line interface.

This chapter discusses the following topics:

- [Getting Started With the ASMCA GUI Tool](#) (page 9-1)
- [Managing Instances and Disk Groups With ASMCA](#) (page 9-2)
- [ASMCA Command-Line Interface](#) (page 9-4)

For information about managing Oracle ACFS and Oracle ADVM with ASMCA, refer to [Managing Oracle ACFS and Oracle ADVM With ASMCA](#) (page 14-1).

Note:

Some ASMCA functionality and menu options are only available in specific environments, such as an Oracle RAC configuration. For details about a specific ASMCA page, access the online help provided with the **Help** button.

9.1 Getting Started With the ASMCA GUI Tool

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 the Oracle ASM instance identified with the Oracle system identifier (SID) set to `+ASM`.

If ASMCA detects an Oracle ASM instance that is not running, then ASMCA prompts whether you want to start the instance.

To start the command-line ASMCA, see "[ASMCA Command-Line Interface](#) (page 9-4)".

Note:

Some ASMCA functionality and menu options are only available in specific environments, such as an Oracle RAC configuration. For details about a specific ASMCA page, access the online help provided with the **Help** button.

9.2 Managing Instances and Disk Groups With ASMCA

This section contains the following topics:

- [Creating Oracle ASM Instances with ASMCA](#) (page 9-2)
- [Managing Disk Groups with ASMCA](#) (page 9-2)

Note:

Some ASMCA functionality and menu options are only available in specific environments, such as an Oracle RAC configuration. For details about a specific ASMCA page, access the online help provided with the **Help** button.

9.2.1 Creating Oracle ASM Instances with ASMCA

Oracle ASM Configuration Assistant enables you to create an Oracle ASM instance if the Oracle ASM instance is not found on the host.

On the **Create ASM** page, you enter the `SYS` password for the administrator of the Oracle ASM instance and the password for the `ASMSNMP` user that can monitor this Oracle ASM instance. `ASMSNMP` is a less privileged user that is primarily used to monitor Oracle ASM instances.

You can select the listener with which to register this Oracle ASM instance. Click **ASM Parameters** to update the initialization parameters for the Oracle ASM instance.

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.

See Also:

- [Initialization Parameters for Oracle ASM Instances](#) (page 3-2) for information about Oracle ASM instance initialization parameters
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for information about installing Oracle ASM
 - *Oracle Database Installation Guide* for information about reviewing database accounts and passwords
-
-

9.2.2 Managing Disk Groups with ASMCA

Oracle ASM Configuration Assistant enables you to configure or create Oracle ASM disk groups with the **Configure ASM Disk Groups** tab.

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.

You can right click a selected disk group in the screen to display the configuration options menu. The menu options enable you to:

- Add disks to and drop disks from the disk group
- Add, drop, and list quota groups.
- Add, drop, and move file groups.
- Edit the disk group attributes
- Manage templates for the disk group
- Create an Oracle ACFS file system on the disk group for database use
- Mount and dismount the disk group
- Drop the disk group
- Drop and clear Oracle ASM Filter Driver labels.

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.

On the **Create Disk Group** 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 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 a quorum failure group is not considered when determining redundancy requirements with respect to user storage. However, a quorum failure group counts when mounting a disk group.

If Oracle ASM Filter Driver has been configured for your system, you can add Oracle ASM Filter Driver labels to disks or clear Oracle ASM Filter Driver labels from disks.

You can choose to show advanced options when creating a disk group. The advanced options section 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.

Note:

The values for **Free (GB)** and **Usable (GB)** space are 0 if the disk group type is `FLEX` or `EXTEND`. The 0 value in this case means no value is given, not a zero amount in for **Free (GB)** and **Usable (GB)** space.

See Also:

- [Storing Oracle Cluster Registry and Voting Files in Oracle ASM Disk Groups](#) (page 4-18) for information about quorum failure groups
- [Administering Oracle ASM Filter Driver](#) (page 20-1) for information about Oracle ASM Filter Driver
- [Disk Group Compatibility](#) (page 4-50) and [Overview of Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility attributes and advancing disk group compatibility
- [Oracle ASM Disk Discovery](#) (page 4-33) and [ASM_DISKSTRING](#) (page 3-10) for information about disk discovery
- [Administering Oracle ASM Disk Groups](#) (page 4-1) for more information about managing disk groups
- *Oracle Database Reference* for descriptions of disk group information displayed in the `V$ASM_DISKGROUP` view

9.3 ASMCA Command-Line Interface

The ASMCA command-line interface provides non-GUI support for configuring Oracle ASM instances, Oracle ASM disk groups, volumes, and Oracle ACFS.

- [Running ASMCA Command-Line](#) (page 9-4)
- [ASMCA Commands for Instances and Disk Groups](#) (page 9-5)

9.3.1 Running ASMCA Command-Line

The syntax for running the ASMCA command-line tool is:

```
asmca -silent command_name [options]
```

[Table 9-1](#) (page 9-4) contains the options available with the `asmca` command-line tool.

Table 9-1 Options for the ASMCA command-line tool

Option	Description
<code>-silent</code>	Specifies to run <code>asmca</code> in command-line mode.

Table 9-1 (Cont.) Options for the ASMCA command-line tool

Option	Description
<i>command_name</i>	<p>Specifies the command to run. The command can be any of the following:</p> <ul style="list-style-type: none"> -configureASM -configureParameter -deleteASM -createDiskGroup -editDiskGroupAttributes -addDisk -createLabel -dropLabel -createFileGroups -dropFileGroups -createQuotaGroups -dropQuotaGroups -createVolume -createACFS -createACFSSnapshot -dropACFSSnapshot <p>See "ASMCA Commands for Instances and Disk Groups" (page 9-5)".</p>
<i>options</i>	<p>Specifies the options to run for a command. The options vary by command. See the individual commands described in "ASMCA Commands for Instances and Disk Groups" (page 9-5)" for the available options.</p> <p>The following list describes the options that are common for multiple commands.</p> <ul style="list-style-type: none"> • -help Displays help for running ASMCA in silent mode. • -asmsnmpPassword <i>asmsnmp_password</i> The Oracle ASM monitor user (ASMSNMP) password. This is optional and only required when creating and migrating an Oracle ASM instance. • -sysAsmPassword <i>sysasm_password</i> 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:
...
```

9.3.2 ASMCA Commands for Instances and Disk Groups

This section describes the commands that can be run with ASMCA command-line to manage Oracle ASM objects.

- [Configure an Oracle ASM Instance](#) (page 9-6)
- [Configure Parameters for a Specified Instance](#) (page 9-7)
- [Delete Oracle ASM Instances](#) (page 9-8)
- [Create a Disk Group](#) (page 9-9)
- [Add a Disk to a Disk Group](#) (page 9-10)
- [Edit Oracle ASM Disk Group Attributes](#) (page 9-11)
- [Create Oracle ASM Filter Driver Labels](#) (page 9-12)
- [Drop Oracle ASM Filter Driver Labels](#) (page 9-12)
- [Create Oracle ASM File Groups](#) (page 9-12)
- [Drop Oracle ASM File Groups](#) (page 9-13)
- [Create Oracle ASM Quota Groups](#) (page 9-13)
- [Drop Oracle ASM Quota Groups](#) (page 9-14)

9.3.2.1 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
        { -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 ] }
```

[Table 9-2](#) (page 9-6) contains the options available with the `-configureASM` command.

Table 9-2 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.

Table 9-2 (Cont.) Options for the `-configureASM` command

Option	Description
<code>-diskGroupName <i>diskgroup</i></code>	Specifies the name of the Oracle ASM disk group to create.
<code>-disk <i>disk_path</i></code>	Specifies the full path of the disk device.
<code>-diskName <i>disk</i></code>	Specifies a user-supplied disk name.
<code>-diskList <i>disk_list</i></code>	Specifies a comma-delimited 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.
<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.
<code>-compatible.rdbms</code>	Specifies the attribute setting.
<code>-compatible.advm</code>	Specifies the attribute setting.

For additional options that are common to multiple commands, see [Table 9-1 \(page 9-4\)](#).

Examples

To configure an Oracle ASM instance:

Example 9-1 Using `asmca -silent -configureASM`

```
asmca -silent
      -configureASM
        -diskString '/devices/disk*'
        -sysAsmPassword my_sysasm_password
        -asmsnmpPassword my_asmsnmp_passwd
```

9.3.2.2 Configure Parameters for a Specified Instance

`-configureParameter` modifies the parameters of the specified target instance.

Syntax

```
asmca -silent
      -configureParameter [-target {ASM|IOS|APX|ALL} ]
      [{-param paramname=paramvalue , ... } ]
      [-sysAsmPassword sysasm_password ]
```

[Table 9-3](#) (page 9-8) contains the options available with the `-configureParameter` command.

Table 9-3 Options for the `-configureParameter` command

Option	Description
<code>-target {ASM IOS APX ALL}</code>	Specifies the target instance of the command. If no target is specified, the default is the Oracle ASM instance.
<code>-param paramname=paramvalue , ...</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 9-1](#) (page 9-4).

Examples

To configure parameters for a specified target instance:

Example 9-2 Using `asmca -silent -configureParameter`

```
asmca -silent
      -configureParameter
        -param
          asm_power_limit=3
      -sysAsmPassword my_sysasm_password
```

9.3.2.3 Delete Oracle ASM Instances

`-deleteASM` deletes all the Oracle ASM instances that are present. 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 9-4](#) (page 9-8) contains the options available with the `-deleteASM` command.

Table 9-4 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 9-1](#) (page 9-4).

Examples

To delete an Oracle ASM instance:

Example 9-3 Using `asmca -silent -deleteASM`

```
asmca -silent
      -deleteASM
      -sysAsmPassword my_sysasm_password
      -force
```

9.3.2.4 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 9-5 (page 9-9) contains the options available with the `-createDiskGroup` command.

Table 9-5 Options for the `-createDiskGroup` 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 to create.
<code>-disk <i>disk_path</i></code>	Specifies the full path of the disk device.
<code>-diskName <i>disk</i></code>	Specifies a user-supplied disk name.
<code>-diskList <i>disk_list</i></code>	Specifies a comma-delimited 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.
<code>-redundancy { HIGH NORMAL EXTERNAL }</code>	Specifies the redundancy setting for the Oracle ASM disk group.

Table 9-5 (Cont.) Options for the `-createDiskGroup` command

Option	Description
<code>-au_size</code>	Specifies the allocation unit size (1,2,4,8,16,32,64)
<code>-compatible.asm</code>	Specifies the attribute setting.
<code>-compatible.rdbms</code>	Specifies the attribute setting.
<code>-compatible.advm</code>	Specifies the attribute setting.

For additional options that are common to multiple commands, see [Table 9-1](#) (page 9-4).

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](#)" (page 4-50).

Examples

To create disk groups in an Oracle ASM instance:

Example 9-4 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.

9.3.2.5 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 9-6](#) (page 9-11) contains the options available with the `-addDisk` command.

Table 9-6 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 to update.
<code>-diskList <i>disk_list</i></code>	Specifies a comma-delimited 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</code> <code>-noforce</code>	Specifies where to force the use of the disk.
<code>-quorum</code> <code>-noquorum</code>	Specifies a quorum or standard failure group.

For additional options that are common to multiple commands, see [Table 9-1](#) (page 9-4).

Examples

To add disks to an existing disk group:

Example 9-5 Using `asmca -silent -addDisk`

```
asmca -silent
      -addDisk
        -diskGroupName mynewdg
          -disk '/devices/disk11'
          -disk '/devices/disk12'
```

Disks added successfully to diskgroup mynewdg.

9.3.2.6 Edit Oracle ASM Disk Group Attributes

`-editDiskGroupAttributes` updates the Oracle ASM disk group attributes.

Syntax and Description

```
-editDiskGroupAttributes
  (-diskGroupName disk_group)
  (-attribute value)
```

The following table contains the options available with the `-editDiskGroupAttributes` command.

Table 9-7 Options for the `-editDiskGroupAttributes` command

Option	Description
<code>-diskGroupName <i>diskgroup</i></code>	Specifies the name of the Oracle ASM disk group to create.

Table 9-7 (Cont.) Options for the `-editDiskGroupAttributes` command

Option	Description
<code>-attribute value</code>	Specifies the name and assigned value of an attribute.

9.3.2.7 Create Oracle ASM Filter Driver Labels

`-createLabel` creates Oracle ASM filter driver labels.

Syntax and Description

```
-createLabel
    -diskList disk_list
    -labelList label_list | -labelPrefix label_prefix
```

The following table contains the options available with the `-createLabel` command.

Table 9-8 Options for the `—createLabel` command

Option	Description
<code>-diskList disk_list</code>	Specifies a comma-delimited list of disk names to be labeled.
<code>-labelList label_list</code>	Specifies a comma-delimited list of label names.
<code>-labelPrefix label_prefix</code>	Specifies the label prefix to be used to create labels. If label list is specified, label prefix is ignored.

9.3.2.8 Drop Oracle ASM Filter Driver Labels

`-dropLabel` drops Oracle ASM filter driver labels.

Syntax and Description

```
-dropLabel
    -labelList label_list
```

The following table contains the options available with the `-dropLabel` command.

Table 9-9 Options for the `—dropLabel` command

Option	Description
<code>-labelList label_list</code>	Specifies a comma-delimited list of label names.

9.3.2.9 Create Oracle ASM File Groups

`-createFileGroups` creates Oracle ASM file groups.

Syntax and Description

```
-createFileGroups
    (-diskGroupName disk_group)
    (-fileGroupNames file_group_name_list)
```

```
(-fileGroupUsages file_group_usages_list)
(-fileGroupUsageIDs file_group_usagesID_list)
```

The following table contains the options available with the `-createFileGroups` command.

Table 9-10 Options for the `-createFileGroups` command

Option	Description
<code>-diskGroupName <i>disk_group</i></code>	Specifies the name of the Oracle ASM disk group.
<code>-fileGroupNames <i>file_group_name_list</i></code>	Specifies a comma-delimited list of file group names to be added to the disk group.
<code>-fileGroupUsages <i>file_group_usages_list</i></code>	Specifies a comma-delimited list of file group usage strings to be processed.
<code>-fileGroupUsageIDs <i>file_group_usagesID_list</i></code>	Specifies a comma-delimited list of file group usage IDs to be processed.

9.3.2.10 Drop Oracle ASM File Groups

`-dropFileGroups` drops Oracle ASM file groups.

Syntax and Description

```
-dropFileGroups
  (-diskGroupName disk_group)
  (-fileGroupNames file_group_name_list)
```

The following table contains the options available with the `-dropFileGroups` command.

Table 9-11 Options for the `-dropFileGroups` command

Option	Description
<code>-diskGroupName <i>disk_group</i></code>	Specifies the name of the Oracle ASM disk group to create.
<code>-fileGroupNames <i>file_group_name_list</i></code>	Specifies a comma-delimited list of file group names to be dropped.

9.3.2.11 Create Oracle ASM Quota Groups

`-createQuotaGroups` creates Oracle ASM quota groups.

Syntax and Description

```
-createQuotaGroups
  (-diskGroupName disk_group)
  (-quotaGroupNames quota_group_name_list)
  (-quotaGroupSizes quota_group_size_list)
```

The following table contains the options available with the `-createQuotaGroups` command.

Table 9-12 Options for the `-createQuotaGroups` command

Option	Description
<code>-diskGroupName <i>disk_group</i></code>	Specifies the name of the Oracle ASM disk group.
<code>-quotaGroupNames <i>quota_group_name_list</i></code>	Specifies a comma-delimited list of quota group names to be added to the disk group.
<code>-quotaGroupSizes <i>quota_group_size_list</i></code>	Specifies a comma-delimited list of quota group sizes to be processed.

9.3.2.12 Drop Oracle ASM Quota Groups

`-dropQuotaGroups` drops Oracle ASM quota groups.

Syntax and Description

```
-dropQuotaGroups
  (-diskGroupName disk_group)
  (-quotaGroupNames quota_group_names_list)
```

The following table contains the options available with the `-dropQuotaGroups` command.

Table 9-13 Options for the `-dropQuotaGroups` command

Option	Description
<code>-diskGroupName <i>disk_group</i></code>	Specifies the name of the Oracle ASM disk group.
<code>-quotaGroupNames <i>quota_group_name_list</i></code>	Specifies a comma-delimited list of quota group names to be dropped.

Managing Oracle ASM with ASMCMD

Oracle Automatic Storage Management (Oracle ASM) Command-Line Utility (ASMCMD) provides commands for managing Oracle ASM.

This chapter describes the Oracle Automatic Storage Management (Oracle ASM) Command-Line Utility (ASMCMD). This chapter contains the following topics:

- [About ASMCMD](#) (page 10-1)
- [ASMCMD Instance Management Commands](#) (page 10-11)
- [ASMCMD Client Cluster Management Commands](#) (page 10-30)
- [ASMCMD Oracle ASM Filter Driver Management Commands](#) (page 10-32)
- [ASMCMD File Management Commands](#) (page 10-42)
- [ASMCMD Disk Group Management Commands](#) (page 10-53)
- [ASMCMD File Group Management Commands](#) (page 10-78)
- [ASMCMD Template Management Commands](#) (page 10-87)
- [ASMCMD File Access Control Commands](#) (page 10-90)

For information about ASMCMD commands to manage volumes, refer to [Managing Oracle ADVM with ASMCMD](#) (page 15-1).

Note:

Oracle does not recommend using identifiers for Oracle Database object names that must be quoted. While the use of quoted identifiers may be valid as names in some command-line tools or in SQL statements, such as `CREATE DISKGROUP "1DATA"`, the names may not be valid when using other tools that manage the object. For more information about naming Oracle Database objects, refer to *Oracle Database SQL Language Reference*.

10.1 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](#) (page 10-2)

- [About Oracle ASM Files, File Names, Directories, and Aliases](#) (page 10-2)
- [Preparing to Run ASMCMD](#) (page 10-6)
- [Running ASMCMD in Interactive Mode](#) (page 10-7)
- [Running ASMCMD in Noninteractive Mode](#) (page 10-10)
- [Getting Help](#) (page 10-11)

10.1.1 Types of ASMCMD Commands

The types of ASMCMD commands are listed in [Table 10-1](#) (page 10-2).

Table 10-1 *Types of ASMCMD commands*

Category	Commands
Oracle ASM instance management	For a list of commands, refer to " ASMCMD Instance Management Commands (page 10-11)"
Oracle ASM client cluster	For a list of commands, refer to ASMCMD Client Cluster Management Commands (page 10-30)
Oracle ASM Filter Driver management	For a list of commands, refer to " ASMCMD Oracle ASM Filter Driver Management Commands (page 10-32)"
Oracle ASM file management	For a list of commands, refer to " ASMCMD File Management Commands (page 10-42)"
Oracle ASM disk group management	For a list of commands, refer to " ASMCMD Disk Group Management Commands (page 10-53)"
Oracle ASM template management	For a list of commands, refer to " ASMCMD Template Management Commands (page 10-87)"
Oracle ASM file access control	For a list of commands, refer to " ASMCMD File Access Control Commands (page 10-90)"
Oracle ADVN volume	Refer to Managing Oracle ADVN with ASMCMD (page 15-1)

10.1.2 About Oracle ASM Files, File Names, 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:

- [System-Generated File Name or Fully Qualified File Name](#) (page 10-3)
- [Directory](#) (page 10-3)
- [Alias](#) (page 10-4)
- [Absolute Path and Relative Path](#) (page 10-4)
- [Wildcard Characters](#) (page 10-5)

10.1.2.1 System-Generated File Name or Fully Qualified File Name

Every file created in Oracle ASM gets a system-generated file name, otherwise known as a fully qualified file name. This is similar to a complete path name in a local file system.

Oracle ASM generates file names according to the following scheme:

```
+diskGroupName/databaseName/fileType/fileTypeTag.fileNumber.incarnation
```

An example of a fully qualified file name is the following:

```
+data/orcl/CONTROLFILE/Current.256.541956473
```

In the previous fully qualified file name, `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. File Names 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 file names, refer to "[Fully Qualified File Name Form](#) (page 5-5)".

10.1.2.2 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 file name represents a hierarchy of directories in which the plus sign (+) represents the root directory. In each disk group, Oracle ASM automatically creates a directory hierarchy that corresponds to the structure of the fully qualified file names 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 file name 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 \(page 10-9\)](#)".

```
ASMCMD [+] > cd data/orcl
ASMCMD [+data/orcl] >
```

ASMCMD retains the case of the directory that you entered.

10.1.2.3 Alias

Aliases are file names that are references or pointers to system-generated file names. 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 file name administration. You can create aliases with the `mkalias` 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 File Name Forms \(page 5-6\)](#)".

10.1.2.4 Absolute Path and Relative Path

When you run an ASMCMD command that accepts a file name 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 file name, otherwise known as the fully qualified file name, 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 file name:

```
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 file name 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
SYS_AUX.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 `".."` rather than 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] >
```

10.1.2.5 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, such as `cd`, `du`, `find`, `ls`, `lsattr`, `lsdg`, `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.

If you are running ASMCMD commands in noninteractive mode, then with most operating systems you must enclose the wildcard characters in quotes. For more information, refer to "[Running ASMCMD in Noninteractive Mode](#) (page 10-10)".

[Example 10-1](#) (page 10-5) illustrates the use of wildcards.

Example 10-1 Using wildcards with ASMCMD commands

```
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
```

10.1.3 Preparing to Run ASMCMD

You should review the information in this topic before running the ASMCMD utility.

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 which contains the Oracle ASM instance that you plan to administer.

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.

- 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.

The default value of the Oracle ASM SID for a single-instance database is `+ASM`. In Oracle 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, 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 `--privilege` option to connect as `SYSDBA`.

With this connection, there is a limited set of operations that can be run.

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.

See Also:

- [Authentication for Accessing Oracle ASM Instances](#) (page 3-27) for more information about authentication for accessing Oracle ASM instances
 - [Specifying the Type of Connection](#) (page 10-9) for more information about specifying connection options
 - [The SYSDBA Privilege for Managing Oracle ASM Components](#) (page 3-31) for more information about permissions and connections
 - *Oracle Database Administrator's Guide* for more information about setting environment variables
 - *Oracle Database Administrator's Guide* for information about operating system authentication
-
-

10.1.4 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.

This section contains the following topics:

- [Syntax for Starting ASMCMD](#) (page 10-8)
- [Running ASMCMD without Any Options](#) (page 10-8)
- [Specifying the Type of Connection](#) (page 10-9)
- [Specifying the Verbose Mode](#) (page 10-9)
- [Including the Current Directory in the ASMCMD Prompt](#) (page 10-9)
- [Displaying the ASMCMD Version Number](#) (page 10-9)

10.1.4.1 Syntax for Starting ASMCMD

The syntax for starting ASMCMD in interactive mode is:

```
asmcmd [-v]
asmcmd [--nocp] [-v {errors|warnings|normal|info|debug}]
          [--privilege connection_type ] [-p]
```

[Table 10-2](#) (page 10-8) summarizes the ASMCMD options.

Table 10-2 ASMCMD interactive options

Option	Description
-v	Displays the version of ASMCMD and then exits.
--nocp	Disables connection pooling. For information about connection pooling, refer to <i>Oracle Database Administrator's Guide</i> .
-v <i>display_level</i>	Displays additional information with some commands to help users diagnose problems and writes to a message file when logging is required. <i>normal</i> is the default if the display level is not specified.
--privilege <i>connection_type</i>	Specifies the privilege to connect as when accessing an Oracle ASM or database instance. Can be either <i>SYSASM</i> or <i>SYSDBA</i> . The default is <i>SYSASM</i> and is used when administering the Oracle ASM instance.
-p	Displays the current directory in the prompt.

ASMCMD can be started with multiple options. For example:

```
asmcmd -p -v info
```

For the majority of the examples in this chapter, ASMCMD was started with the `-p` option.

10.1.4.2 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.

10.1.4.3 Specifying the Type of Connection

You can specify the `--privilege` option to choose the type of connection, either `SYSASM` or `SYSDBA`. The default value is `SYSASM` and is used when administering an Oracle ASM instance. Connect as `SYSDBA` when connecting to the database instance.

For example:

```
$ asmcmd --privilege sysasm
```

10.1.4.4 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 debug
```

When there is a message that must be logged during `asmcmd` operations, the message is written to the `alert.log` file or the `trace.trc` file.

The `alert.log` file is created or updated when `asmcmd` is started with the `-v errors`, `-v warning`, or `-v normal` display level. The `trace.trc` file is created or updated when `asmcmd` is started with the `-v info` or `-v debug` display level.

By default, the `alert.log` file is created in the `$ORACLE_HOME/log/diag/asmcmd/user_username/hostname/alert/` directory. By default, the `trace.trc` file is created in the `$ORACLE_HOME/log/diag/asmcmd/user_username/hostname/trace/` directory.

Under certain circumstances, `$ORACLE_BASE` and `$ORACLE_HOME` can be set to override the default locations of the `alert.log` and `trace.trc` files.

10.1.4.5 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] >
```

10.1.4.6 Displaying the ASMCMD Version Number

You can specify the `-V` option when starting `asmcmd` to display the `asmcmd` version number. After displaying the version number, `asmcmd` immediately exits.

For example:

```
$ asmcmd -V
asmcmd version 11.2.0.3.0
$
```

10.1.5 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 when invoking ASMCMD. ASMCMD runs the command, generates output if any, and then exits. The noninteractive mode is especially useful for running scripts.

On most operating systems, wildcard characters must be enclosed in quotes when running ASMCMD commands in noninteractive mode.

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
```

In noninteractive mode, ASMCMD returns the codes listed in [Table 10-3](#) (page 10-10).

Table 10-3 ASMCMD return codes

Type	Return Code	Description
General	0	No issues for the command.
General	1	Internal error.
General	255 or -1	External user error as returned by a command.
lsdsk	0	Results include member disks only.
lsdsk	1	Results include candidate disks only.
lsdsk	2	Results include both member and candidate disks.

[Example 10-2](#) (page 10-10) shows how to run ASMCMD in the noninteractive mode. The first command runs the `ls` command to list the disk groups for the Oracle ASM instance. The second command redirects the output of the `lsod` command to the `my_lsod_test` file. The third command runs `lsdsk` using a pattern with a wildcard character to list specific disks in the `data` disk group. Note that wildcard characters must be enclosed in quotes when running ASMCMD in noninteractive mode on most operating systems.

[Example 10-3](#) (page 10-11) raises an error with an incorrect disk group name and the return code equal (255) is displayed.

[Example 10-4](#) (page 10-11) shows an example of ASMCMD commands run in noninteractive mode inside a script.

Example 10-2 Running ASMCMD commands in noninteractive mode

```
$ asmcmd ls -l
State   Type   Rebal  Name
MOUNTED NORMAL N      DATA/
MOUNTED NORMAL N      FRA/

$ asmcmd lsod --suppressheader -G data > my_lsod_test

$ asmcmd lsdsk -G data '/devices/diska*'
Path
/devices/diskal
```



```
/devices/diska2
/devices/diska3
```

Example 10-3 Displaying a return code when running ASMCMD in noninteractive mode

```
$ asmcmd ls -l dat
ASMCMD-08001: diskgroup 'dat' does not exist or is not mounted

$ echo $?
255
```

Example 10-4 Running ASMCMD commands in a script

```
#!/bin/sh
for ((i = 1; i <=3; i++))
do
    asmcmd lsdsk -G data '/devices/diska'$i
done

$ ./asmcmd_test_script
Path
/devices/diska1
Path
/devices/diska2
Path
/devices/diska3
```

10.1.6 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 10-5 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 file.
```

10.2 ASMCMD Instance Management Commands

This topic provides a summary of the ASMCMD 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.

See Also:

Oracle Grid Infrastructure Installation and Upgrade Guide for your operating system for information about installing and configuring Oracle Grid Infrastructure

[Table 10-4](#) (page 10-12) lists the Oracle ASM instance management commands with brief descriptions.

Table 10-4 Summary of ASMCMD instance management commands

Command	Description
dsget (page 10-13)	Retrieves the discovery diskstring value that is used by the Oracle ASM instance and its clients.
dsset (page 10-13)	Sets the disk discovery diskstring value that is used by the Oracle ASM instance and its clients.
lsct (page 10-14)	Lists information about current Oracle ASM clients.
lsop (page 10-15)	Lists the current operations on a disk group or Oracle ASM instance.
lspwusr (page 10-15)	Lists the users from an Oracle ASM password file.
orapwusr (page 10-16)	Adds, drops, or changes an Oracle ASM password user.
pwcopu (page 10-17)	Copies a password file to the specified location.
pwcreate (page 10-18)	Creates a password file at the specified location.
pwdelete (page 10-19)	Deletes a password file at the specified location.
pwget (page 10-19)	Returns the location of the password file.
pwmov (page 10-20)	Moves the location of the password file.
pwset (page 10-21)	Sets the location of the password file.
showclustermode (page 10-22)	Displays the current mode of the Oracle ASM cluster.
showclusterstate (page 10-22)	Displays the current state of the cluster.
showpatches (page 10-22)	Lists the patches applied to the Oracle Grid Infrastructure home.
showversion (page 10-23)	Displays the Oracle ASM cluster release and software patch levels.
shutdown (page 10-23)	Shuts down an instance.
spbackup (page 10-24)	Backs up an Oracle ASM SPFILE.
spcopy (page 10-25)	Copies an Oracle ASM SPFILE.
spget (page 10-27)	Retrieves the location of the Oracle ASM SPFILE.
spmov (page 10-27)	Moves an Oracle ASM SPFILE.
spset (page 10-28)	Sets the location of the Oracle ASM SPFILE.
startup (page 10-29)	Starts up an instance.

10.2.1 dsget

Purpose

Retrieves the discovery diskstring value that is used by the Oracle ASM instance and its clients.

Syntax and Description

```
dsget [ --normal | --parameter | --profile [-f] ]
```

The syntax options for the `dsget` command are described in [Table 10-5](#) (page 10-13).

Table 10-5 Options for the `dsget` command

Option	Description
<code>--normal</code>	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.
<code>--parameter</code>	Retrieves the <code>ASM_DISKSTRING</code> parameter setting of the Oracle ASM instance.
<code>--profile [-f]</code>	Retrieves the discovery string from the GPnP profile. If <code>-f</code> is specified with <code>--profile</code> , <code>dsget</code> retrieves the discovery string from the local GPnP profile.

Example

The following example uses `dsget` to retrieve the current discovery diskstring value from the GPnP profile and the `ASM_DISKSTRING` parameter.

Example 10-6 Using the ASMCMD `dsget` command

```
ASMCMD [+] > dsget
profile: /devices1/disk*
parameter: /devices2/disk*
```

10.2.2 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 [-f] ] diskstring
```

The syntax options for the `dsset` command are described in [Table 10-6](#) (page 10-14).

Table 10-6 Options for the dsset command

Option	Description
<code>--normal</code>	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.
<code>--parameter</code>	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.
<code>--profile [-f]</code>	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 <code>-f</code> is specified with <code>--profile</code> , 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 for recovery. The command fails if the Oracle Clusterware stack is running.
<i>diskstring</i>	Specifies the value for the discovery diskstring.

For information about disk discovery and the discovery diskstring, see "[Oracle ASM Disk Discovery](#) (page 4-33)".

Example

The following example uses `dsset` to set the current value of the discovery diskstring in the GPnP profile.

Example 10-7 Using the ASMCMD dsset command

```
ASMCMD [+] > dsset /devices1/disk*,/devices2/disk*
```

10.2.3 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 [--suppressheader] [-g] diskgroup
```

[Table 10-7](#) (page 10-15) lists the options for the `lsct` command.

Table 10-7 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.
--suppressheader	Suppresses column headings.
diskgroup	Specifies the disk group.

Client information is only listed for the specified disk group.

Example

The following example displays information about the clients that are accessing the data disk group.

Example 10-8 Using the ASMCMD lsct command

```
ASMCMD [+] > lsct data
DB_Name Status Software_Version Compatible_version Instance_Name Disk_Group
+ASM CONNECTED 12.1.0.0.1 12.1.0.0.1 +ASM DATA
asmvol CONNECTED 12.1.0.0.1 12.1.0.0.1 +ASM DATA
orcl CONNECTED 12.1.0.0.1 12.0.0.0.0 orcl DATA
```

10.2.4 lsop

Purpose

Lists the current operations on a disk group in an 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 10-9 Using the ASMCMD lsop command

```
ASMCMD [+] > lsop
Group_Name Dsk_Num State Power
DATA REBAL WAIT 2

ASMCMD [+] > lsop
Group_Name Dsk_Num State Power
FRA REBAL REAP 3
```

10.2.5 lspwusr

Purpose

List the users from the local Oracle ASM password file.

Syntax and Description

```
lspwusr [--suppressheader]
```

Table 10-8 (page 10-16) lists the options for the `lspwusr` command.

Table 10-8 Options for the `lspwusr` command

Option	Description
<code>--suppressheader</code>	Suppresses column headers from the output.

Examples

The following is an example of the `lspwusr` example. The example lists the current users in the local Oracle ASM password file.

Example 10-10 Using the ASMCMD `lspwusr` command

```
ASMCMD [+] > lspwusr
Username sysdba sysoper sysasm
        SYS   TRUE   TRUE   TRUE
ASMSNMP  TRUE   FALSE  FALSE
```

10.2.6 orapwusr

Purpose

Add, drop, or modify an Oracle ASM password file user.

Syntax and Description

```
orapwusr { --add | --modify | --delete | grant {sysasm|sysdba|sysoper} |
          --revoke {sysasm|sysdba|sysoper} } user
```

Table 10-9 (page 10-16) lists the options for the `orapwusr` command.

Table 10-9 Options for the `orapwusr` command

Option	Description
<code>--add</code>	Adds a user to the password file. Also prompts for a password.
<code>--modify</code>	Changes the password for the specified user.
<code>--delete</code>	Drops a user from the password file.
<code>--grant role</code>	Sets the role for the user. The options are <code>sysasm</code> , <code>sysdba</code> , and <code>sysoper</code> .
<code>--revoke role</code>	Removes the role for the user. The options are <code>sysasm</code> , <code>sysdba</code> , and <code>sysoper</code> .
<code>user</code>	Name of the user to add, drop, or modify.

`orapwusr` attempts to update passwords on all nodes in a cluster. This 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.

Example 10-11 Using the ASMCMD orapwusr command

```
ASMCMD [+] > orapwusr --add hruser
```

10.2.7 pwcopu

Purpose

Copies an Oracle ASM or database instance password file to the specified location.

Syntax and Description

```
pwcopu [--asm | --dbuniquename string] source destination [-f]
```

Table 10-10 (page 10-17) lists the options for the `pwcopu` command.

Table 10-10 Options for the pwcopu command

Option	Description
<code>--asm</code>	The <code>--asm</code> option associates the password file with the Oracle ASM instance that ASMCMD is logged into.
<code>--dbuniquename <i>string</i></code>	The <code>--dbuniquename <i>string</i></code> option identifies the database unique name associated with the password file.
<i>source</i>	The <i>source</i> value identifies the location where the existing password file is located.
<i>destination</i>	The <i>destination</i> value identifies the location where you want to copy the password file.
<code>-f</code>	The <code>-f</code> option enables the password file to be copied without any checks.

`pwcopu` copies a password file from one disk group to another, from the operating system to a disk group, or from a disk group to the operating system. If the `-f` option is specified, then the password file can be copied to the same disk group.

Either `--asm` or `--dbuniquename` is required to identify a CRSD resource. When either `--asm` or `--dbuniquename` is included in the `pwcopu` command, the target file is set to the current password file.

The `compatible.asm` disk group attribute must be set to 12.1 or higher for the disk group where the password is to be copied.

The `SYSASM` or `SYSDBA` privilege is required to manage the Oracle ASM and database password files.

Example

The following example copies an Oracle ASM password file in one disk group to a different disk group. Because the `pwcopu` command includes the `--asm` option, the target file (`+fra/orapwasm_new`) is set to the current password file.

Example 10-12 Using the ASMCMD pwcopy command

```
ASMCMD [+] > pwcopy --asm +DATA/orapwasm +FRA/orapwasm_new
copying +DATA/orapwasm -> +FRA/orapwasm_new
```

10.2.8 pwcreate

Purpose

Note:

Do not use `pwcreate` to create a new Oracle ASM password file when Oracle ASM is configured in a cluster. If the Oracle ASM password file is inaccessible in a cluster configuration, you must restore the password file using an available backup rather than create a new file. For articles about recreating a shared Oracle ASM password file in a cluster, such as document 1929673.1, refer to articles at [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com).

Creates an Oracle ASM or database instance password file at the specified location.

Syntax and Description

```
pwcreate { --asm | --dbuniquename string } file_path sys_password
```

[Table 10-11](#) (page 10-18) lists the options for the `pwcreate` command.

Table 10-11 Options for the pwcreate command

Option	Description
<code>--asm</code>	The <code>--asm</code> option associates the password file with the Oracle ASM instance that ASMCMD is logged into.
<code>--dbuniquename <i>string</i></code>	The <code>--dbuniquename <i>string</i></code> option identifies the database unique name associated with the password file.
<code><i>file_path</i></code>	The <code><i>file_path</i></code> value identifies the location where the password file is created.
<code><i>sys_password</i></code>	The <code><i>sys_password</i></code> value identifies the initial SYS password.

`pwcreate` creates a password file in the disk group specified by `file_path` with an initial SYS password identified by `sys_password`.

Either `--asm` or `--dbuniquename` is required. The `compatible.asm` disk group attribute must be set to 12.1 or higher for the disk group where the password is to be located.

The SYSASM or SYSDBA privilege is required to manage the Oracle ASM and database password files.

For additional information, refer to [Managing a Shared Password File in a Disk Group](#) (page 3-33).

Example

The following example creates an Oracle ASM password file in an Oracle ASM disk group.

Example 10-13 Using the ASMCMD pwcreate command

```
ASMCMD [+] > pwcreate --asm '+DATA/orapwasm' 'welcome'
```

10.2.9 pwdelete

Purpose

Deletes an Oracle ASM or database instance password file.

Syntax and Description

```
pwdelete { --asm | --dbuniquename string | file_path }
```

Table 10-12 (page 10-19) lists the options for the pwdelete command.

Table 10-12 Options for the pwdelete command

Option	Description
--asm	The --asm option associates the password file with the Oracle ASM instance that ASMCMD is logged into.
--dbuniquename <i>string</i>	The --dbuniquename <i>string</i> option identifies the database unique name associated with the password file.
<i>file_path</i>	The <i>file_path</i> value identifies the location where the password file is located.

pwdelete deletes the specified password file. Either --asm or --dbuniquename is required to identify a CRSD resource and to remove the password location from the CRSD resource.

The SYSASM or SYSDBA privilege is required to manage the Oracle ASM and database password files.

Example

The following example deletes the specified password file from a disk group.

Example 10-14 Using the ASMCMD pwdelete command

```
ASMCMD [+] > pwdelete +FRA/orapwasm_bak
```

10.2.10 pwget

Purpose

Returns the location of the password file for the Oracle ASM or database instance.

Syntax and Description

```
pwget { --asm | --dbuniquename string }
```

Table 10-13 (page 10-20) lists the options for the pwget command.

Table 10-13 Options for the pwget command

Option	Description
<code>--asm</code>	The <code>--asm</code> option associates the password file with the Oracle ASM instance that ASMCMD is logged into.
<code>--dbunique <i>string</i></code>	The <code>--dbunique <i>string</i></code> option identifies the database unique name associated with the password file.

`pwget` returns the location of the password file for the Oracle ASM instance identified by `--asm` or the database instance identified by `--dbunique`.

The `SYSASM` or `SYSDBA` privilege is required to manage the Oracle ASM and database password files.

Example

The following example returns the location of the Oracle ASM password file.

Example 10-15 Using the ASMCMD pwget command

```
ASMCMD [+] > pwget --asm
+DATA/orapwasm
```

10.2.11 pwmove

Purpose

Moves an Oracle ASM or database instance password file to the specified location.

Syntax and Description

```
pwmove { --asm | --dbunique string } source destination [-f]
```

[Table 10-14](#) (page 10-20) lists the options for the `pwmove` command.

Table 10-14 Options for the pwmove command

Option	Description
<code>--asm</code>	The <code>--asm</code> option associates the password file with the Oracle ASM instance that ASMCMD is logged into.
<code>--dbunique <i>string</i></code>	The <code>--dbunique <i>string</i></code> option identifies the database unique name associated with the password file.
<i>source</i>	The <i>source</i> value identifies the location where the existing password file is located.
<i>destination</i>	The <i>destination</i> value identifies the location where you want to move the password file.
<code>-f</code>	The <code>-f</code> option clears the password file for any associated resource and the new file is registered.

`pwmove` moves a password file from one disk group to another, from the operating system to a disk group, or from a disk group to the operating system. If the `-f` option is specified, then a password file can be moved to a file in the same disk group.

Either `--asm` or `--dbunique` is required to identify a CRSD resource.

The `compatible.asm` disk group attribute must be set to 12.1 or higher for the disk group where the password is to be moved.

The `SYSASM` or `SYSDBA` privilege is required to manage the Oracle ASM and database password files.

Example

The following example moves a password file from one disk group to another disk group.

Example 10-16 Using the ASMCMD pwmove command

```
ASMCMD [+] > pwmove --asm +FRA/orapwasm_bak +DATA/orapwasm
moving +FRA/orapwasm_bak -> +DATA/orapwasm
```

10.2.12 pwset

Purpose

Sets the location of the password file for an Oracle ASM or database instance.

Syntax and Description

```
pwset { --asm | --dbunique string } file_path
```

[Table 10-15](#) (page 10-21) lists the options for the `pwset` command.

Table 10-15 Options for the pwset command

Option	Description
<code>--asm</code>	The <code>--asm</code> option associates the password file with the Oracle ASM instance that ASMCMD is logged into.
<code>--dbunique <i>string</i></code>	The <code>--dbunique <i>string</i></code> option identifies the database unique name associated with the password file.
<code><i>file_path</i></code>	The <code><i>file_path</i></code> value identifies the location where the password file is located.

`pwset` sets the location of the password file for an Oracle ASM or database instance to the value specified by `file_path`. Either `--dbunique` or `--asm` is required to identify a CRSD resource.

The `SYSASM` or `SYSDBA` privilege is required to manage the Oracle ASM and database password files.

Example

The following example sets the location of the Oracle ASM password file in a disk group.

Example 10-17 Using the ASMCMD pwset command

```
ASMCMD [+] > pwset --asm +DATA/orapwasm
```

10.2.13 showclustermode

Purpose

Displays the current mode of the Oracle ASM cluster.

Syntax and Description

```
showclustermode
```

showclustermode displays the mode that the Oracle ASM cluster is currently in. The possible return values are **ASM cluster : Flex mode enabled** or **ASM cluster : Flex mode disabled**.

Example

The following example shows the use of the showclustermode command.

Example 10-18 Using the ASMCMD showclustermode command

```
ASMCMD [+] > showclustermode
ASM cluster : Flex mode disabled
```

10.2.14 showclusterstate

Purpose

Displays the current state of the cluster.

Syntax and Description

```
showclusterstate
```

showclusterstate displays the state that the Oracle ASM cluster is currently in. The possible values returned are **normal**, **in-upgrade**, or **in-rollingpatch** mode.

Example

This example shows the use of the showclusterstate command.

Example 10-19 Using the ASMCMD showclusterstate command

```
ASMCMD [+] > showclusterstate
```

10.2.15 showpatches

Purpose

Lists the patches applied on the Oracle Grid Infrastructure home.

Syntax and Description

```
showpatches
```

showpatches lists the patches that have been applied to the Oracle Grid Infrastructure home.

Example

This example shows the use of the showpatches command.

Example 10-20 Using the ASMCMD showpatches command

```
ASMCMD [+] > showpatches
-----
List of Patches
-----
```

10.2.16 showversion

Purpose

Displays the patch levels of the Oracle ASM cluster release and software.

Syntax and Description

```
showversion [--releasepatch] [--softwarepatch]
```

Table 10-16 (page 10-23) lists the options for the showversion command.

Table 10-16 Options for the showversion command

Option	Description
--releasepatch	Displays the cluster level patch version.
--softwarepatch	Displays the patch level on the local node.

showversion displays the Oracle ASM cluster release and software patch levels. The release and software patch options may differ depending whether the Oracle ASM or Oracle Grid Infrastructure home is being patched.

Example

This example shows the use of the showversion command.

Example 10-21 Using the ASMCMD showversion command

```
ASMCMD [+] > showversion --softwarepatch
ASM version      : 12.1.0.1.0
Software patchlevel : 0

ASMCMD [+] > showversion --releasepatch
ASM version      : 12.1.0.1.0
```

10.2.17 shutdown

Purpose

Shuts down an instance.

Syntax and Description

```
shutdown [--target target_instance] [--normal | --abort|--immediate ]
```

[Table 10-17](#) (page 10-24) lists the options for the shutdown command.

Table 10-17 Options for the shutdown command

Option	Description
--target <i>target_instance</i>	Specify the target instance. The value of <i>target_instance</i> could be either Oracle ASM (ASM), IO Server (IOS), or Oracle ASM proxy (APX) instance.
--normal	Shut down normal. This is the default action.
--abort	Shut down aborting all existing operations.
--immediate	Shut down immediately.

The default target instance is determined by the ORACLE_SID environmental variable. 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 "[About Shutting Down an Oracle ASM Instance](#) (page 3-22)".

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 10-22 Using the ASMCMD shutdown command

```
ASMCMD [+] > shutdown --target ASM --normal
ASMCMD [+] > shutdown --target ASM --immediate
ASMCMD [+] > shutdown --target ASM --abort
```

10.2.18 spbackup

Purpose

Backs up an Oracle ASM SPFILE to a backup file.

Syntax and Description

```
spbackup source destination
```

[Table 10-18](#) (page 10-25) lists the options for the spbackup command.

Table 10-18 Options for the `spbackup` command

Option	Description
<code>source</code>	Specifies the source file name.
<code>destination</code>	Specifies the destination file.

`spbackup` should be used when you want to make single or multiple backups of an SPFILE in the same or a different disk group without creating an SPFILE in the target disk group.

Note the following about the use of `spbackup`:

- `spbackup` can back up an Oracle ASM SPFILE from a disk group to a disk group or to an operating system file.
- `spbackup` can back up an Oracle ASM SPFILE from an operating system file to a disk group.
- `spbackup` can back up an Oracle ASM SPFILE when the SPFILE is being used by an open Oracle ASM instance.
- `spbackup` can make multiple backups of an Oracle ASM SPFILE in the same disk group.

`spbackup` does not affect the GPnP profile. The backup file that is created is not a special file type and is not identified as an SPFILE. This backup file cannot be copied with `spcopy`. To copy this backup file to and from a disk group, use the ASMCMD `cp` command.

To make a copy of a backup file in a disk group that is identified as an SPFILE file:

1. Use the ASMCMD `cp` command to copy the backup file from the disk group to an operating system file. See "[cp](#) (page 10-44)".
2. Use the ASMCMD `spcopy` command to copy the operating system file to a disk group. See "[spcopy](#) (page 10-25)".

Example

The following are examples of the `spbackup` command. The first example backs up the SPFILE in the `data` disk group. The second example backs up the SPFILE from the `data` disk group to the `fra` disk group.

Example 10-23 Using the ASMCMD `spbackup` command

```

ASMCMD> spbackup +DATA/asm/asmparameterfile/registry.253.721810181
                +DATA/spfileBackASM.bak

ASMCMD> spbackup +DATA/asm/asmparameterfile/registry.253.721810181
                +FRA/spfileBackASM.bak
    
```

10.2.19 `spcopy`

Purpose

Copies an Oracle ASM SPFILE from the source location to an SPFILE in the destination location.

Syntax and Description

`spcopy [-u] source destination`

Table 10-19 (page 10-26) lists the options for the `spcopy` command.

Table 10-19 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.
- `spcopy` can copy an Oracle ASM SPFILE when the SPFILE is being used by an open Oracle ASM instance.
- `spcopy` cannot make multiple copies of an Oracle ASM SPFILE in the same disk group. You can use `spbackup` for that purpose.

To update the GPnP profile, include the `-u` option with `spcopy`. You can also use `spset` to update the GPnP profile if `spcopy` is run without the `-u` option.

After copying the SPFILE and updating the GPnP profile, you must restart the instance with the SPFILE in the new location 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 copy an Oracle ASM SPFILE into a disk group using `spcopy`, the `COMPATIBLE .ASM` attribute must be set to 11.2 or greater in the target disk group.

Example

The following are examples of the `spcopy` command. The first example copies the Oracle ASM SPFILE from the `data` disk group to the `fra` disk group. The second example copies the Oracle ASM SPFILE from the `data` disk group to an operating system location. The third example copies an Oracle ASM SPFILE from an operating system location to the `data` disk group and updates the GPnP profile with the `-u` option.

Example 10-24 Using the ASMCMD `spcopy` command

```
ASMCMD> spcopy +DATA/asm/asmparameterfile/registry.253.721810181
+FRA/spfileCopyASM.ora
```

```
ASMCMD> spcopy +DATA/asm/asmparameterfile/registry.253.721810181
$ORACLE_HOME/dbs/spfileCopyASM.ora
```

```
ASMCMD> spcopy -u /oracle/product/11.2.0/grid/dbs/spfileTestASM.ora
+DATA/ASM/spfileCopyASM.ora
```


See Also:

- [spset](#) (page 10-28) for information about using the ASMCMD `spset` command
- [About Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File](#) (page 3-4) for information about copying and moving an Oracle ASM instance initialization parameter file after upgrading
- *Oracle Database SQL Language Reference* for information about creating a server parameter file with the `CREATE SPFILE` SQL statement

10.2.20 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 Oracle 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 10-25 Using the ASMCMD `spget` command

```
ASMCMD [+] > spget
+DATA/ASM/ASMPARAMETERFILE/registry.253.813507611
```

10.2.21 spmove

Purpose

Moves an Oracle ASM SPFILE from source to destination and automatically updates the GPnP profile.

Syntax and Description

```
spmove source destination
```

[Table 10-20](#) (page 10-28) lists the options for the `spmove` command.

Table 10-20 Options for the *spmove* command

Option	Description
<i>source</i>	Specifies the source file.
<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 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 "[About Backing Up, Copying, and Moving an Oracle ASM Initialization Parameter File](#) (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 an Oracle ASM SPFILE from the `data` disk group to an operating system location. The second example moves an SPFILE from an operating system location to the `data` disk group.

Example 10-26 Using the ASMCMD *spmove* command

```
ASMCMD> spmove +DATA/spfileASM.ora
           /oracle/product/11.2.0/grid/dbs/spfileMoveASM.ora

ASMCMD> spmove /oracle/product/11.2.0/grid/dbs/spfile+ASM.ora
           +DATA/ASM/spfileMoveASM.ora
```

10.2.22 spset

Purpose

Sets the location of the Oracle ASM SPFILE in the Grid Plug and Play (GPnP) profile.

Syntax and Description

```
spset location
```

[Table 10-21](#) (page 10-28) lists the options for the *spset* command.

Table 10-21 Options for the *spset* command

Option	Description
<i>location</i>	Specifies the location of the Oracle ASM SPFILE. The location is the full path to the 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 10-27 Using the ASMCMD spset command

```
ASMCMD> spset +DATA/asm/asmparameterfile/asmspfile.ora
```

10.2.23 startup

Purpose

Starts up the default instance.

Syntax and Description

```
startup [--mount] [--nomount] [--restrict] [ --pfile pfile ]
```

[Table 10-22](#) (page 10-29) lists the options for the `startup` command.

Table 10-22 Options for the startup command

Option	Description
<code>--mount</code>	Specifies a mount operation. This is the default action.
<code>--nomount</code>	Specifies no mount operation.
<code>--restrict</code>	Specifies restricted mode.
<code>--pfile <i>pfile</i></code>	Oracle ASM initialization parameter file.

This command starts up an instance specified by the `ORACLE_SID` environmental variable. The variable could be set to the Oracle ASM, IO Server, or APX proxy instance.

The default action for an Oracle ASM instance is a startup that mounts disk groups and enables Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes.

Example

The following is an example of the `startup` command that starts the Oracle ASM instance (`ORACLE_SID=+ASM`) without mounting disk groups and uses the `asm_init.ora` initialization parameter file.

Example 10-28 Using the ASMCMD startup command

```
ASMCMD> startup --nomount --pfile asm_init.ora
```

See Also:

- [About Mounting Disk Groups at Startup](#) (page 3-21) for information about disk groups that are mounted at startup time
- [About Starting Up an Oracle ASM Instance](#) (page 3-18) for more information about starting up an Oracle ASM instance

10.3 ASMCMD Client Cluster Management Commands

This topic provides a summary of the ASMCMD client cluster management commands that are used with Oracle ASM remote instances.

Unless otherwise specified, the ASMCMD cluster management commands should be executed in the Oracle ASM-CSS cluster.

For information about Oracle Flex ASM, refer to "[Managing Oracle Flex ASM](#) (page 19-1)".

[Table 10-23](#) (page 10-30) lists the ASMCMD client cluster management commands with brief descriptions.

Table 10-23 Summary of ASMCMD client cluster management commands

Command	Description
mkcc (page 10-31)	Creates a configuration for a member cluster.
lsc (page 10-30)	Lists the configured member clusters.
rmcc (page 10-32)	Removes the configuration for a member cluster.

10.3.1 lsc

Purpose

Lists the configured member clusters.

Syntax and Description

```
lsc [--suppressheader] [-1] [cluster_name]
```

The following table lists the options for the `lsc` command.

Table 10-24 Options for the lsc command

Option	Description
<code>--suppressheader</code>	Suppresses column headings from the output.
<code>-1</code>	Displays all details.
<code>cluster_name</code>	Specifies the name of the client cluster.

Examples

The following example lists the configured client clusters.

Example 10-29 Using the ASMCMD lsc command

```
ASMCMD [+] > lsc
-----
          NAME          VERSION GUID
-----
ClientCluster1      12.2.0.0.0 78015b1e78756f88bf119fd8e5146df4
Clientcluster2      12.2.0.0.0 3de3bbd75e9a7f17bfe027b477b5bcae
-----
```

```

ASMCMD [+] > lsc -l
-----
          NAME          VERSION GUID          ASM GIMR TFA RHP
=====
ClientCluster1    12.2.0.0.0 78015ble78756f88bf119fd8e5146df4 YES  YES YES YES
Clientcluster2    12.2.0.0.0 3de3bbd75e9a7f17bfe027b477b5bcae YES  YES YES NO
=====

ASMCMD [+] > lsc ClientCluster1 --suppreshheader
ClientCluster1    12.2.0.0.0 78015ble78756f88bf119fd8e5146df4

ASMCMD [+] > lsc ClientCluster1 --suppreshheader -l
ClientCluster1    12.2.0.0.0 78015ble78756f88bf119fd8e5146df4 ASM,GIMR,TFA,RHP

```

10.3.2 mkcc

Purpose

Creates the configuration for a member cluster.

Syntax and Description

```

mkcc cluster_name wrap [--direct] [--version client_cluster_version] [--asm] [--gimr] [--rhp] [--tfa]

```

The following table lists the options for the mkcc command.

Table 10-25 Options for the mkcc command

Option	Description
<i>cluster_name</i>	Specifies the name of the client cluster.
<i>wrap</i>	Specifies the full path of the XML file to export the credentials.
--direct	Specifies direct storage access. If not specified, the storage access is indirect.
--version <i>client_cluster_version</i>	Specifies the five digit client cluster version if it is different than the storage server version. The storage server version is used if not specified.
--asm	Generate credentials and configuration for an Oracle ASM client cluster.
--gimr	Generate credentials and configuration for an Oracle GIMR client cluster.
--rhp	Generate credentials and configuration for an Oracle RHP client cluster.
--tfa	Generate credentials and configuration for an Oracle TFA client cluster.

The client cluster name must conform to the same rules as a cluster name. The name must be unique across your enterprise, must be at least one character long and no more than 15 characters in length, must be alphanumeric, cannot begin with a numeral, and may contain hyphens (-). Underscore characters (_) are not allowed.

If no component options are specified, the default is to configure all of the components and store their credentials in the same manifest file.

GNS credentials are not generated by the `asmcmd mkcc` command. To configure the member cluster as a GNS Client, run the following command as a privileged user to export the GNS instance client data configuration to the cluster manifest file:

```
srvctl export gns -clientdata wrap -role CLIENT
```

Examples

The following example creates credentials for a client cluster.

Example 10-30 Using the ASMCMD mkcc command

```
ASMCMD [+] > mkcc clientcluster1 /tmp/client_cluster1.xml --version 12.2.0.0.0
```

See Also:

Oracle Grid Infrastructure Installation and Upgrade Guide for information about cluster name requirements

10.3.3 rmcc

Purpose

Removes the configuration for a member cluster.

Syntax and Description

```
rmcc cluster_name [-f]
```

The following table lists the options for the `rmcc` command.

Table 10-26 Options for the rmcc command

Option	Description
<code>cluster_name</code>	Specifies the name of the client cluster.
<code>-f</code>	Specifies to force the removal of client cluster credentials.

Examples

The following example removes the configuration for a client cluster.

Example 10-31 Using the ASMCMD rmcc command

```
ASMCMD [+] > rmcc clientcluster1
```

10.4 ASMCMD Oracle ASM Filter Driver Management Commands

This topic provides a summary of the ASMCMD Oracle ASM filter driver (Oracle ASMFDF) management commands.

See Also:

Oracle Grid Infrastructure Installation Guide for your operating system for information about installing and configuring Oracle Grid Infrastructure

Table 10-27 (page 10-33) lists the Oracle ASM Filter Driver management commands with brief descriptions.

Table 10-27 Summary of ASMCMD Oracle ASM Filter Driver management commands

Command	Description
afd_configure (page 10-33)	Configures Oracle ASM Filter Driver.
afd_deconfigure (page 10-34)	Deconfigures Oracle ASM Filter Driver.
afd_di (page 10-35)	Manipulates the data integrity state in the Oracle ASM Filter Driver.
afd_dsget (page 10-36)	Retrieves the Oracle ASM Filter Driver diskstring value.
afd_dsset (page 10-36)	Sets the Oracle ASM Filter Driver diskstring value.
afd_filter (page 10-36)	Sets the Oracle ASM Filter Driver mode on a given disk path.
afd_label (page 10-37)	Sets an Oracle ASM Filter Driver label to a disk.
afd_lsdsdsk (page 10-38)	Lists Oracle ASM Filter Driver disks.
afd_lslbl (page 10-38)	Lists the disks that contain Oracle ASM Filter Driver labels.
afd_refresh (page 10-39)	Refreshes Oracle ASM Filter Driver disks using the AFD discovery string.
afd_scan (page 10-40)	Scans for Oracle ASM Filter Driver disks.
afd_state (page 10-40)	Returns the state of Oracle ASM Filter Driver.
afd_unlabel (page 10-41)	Clears an existing Oracle ASM Filter Driver label.

10.4.1 afd_configure

Purpose

Configures Oracle ASM Filter Driver.

Syntax and Description

```
# asmcmd afd_configure [-d | -e] [-f]
```

The following table lists the options for the `afd_configure` command.

Table 10-28 Options for the afd_configure command

Option	Description
-d	Disables Oracle ASMFD filtering mode.
-e	Enables Oracle ASMFD filtering mode.

Table 10-28 (Cont.) Options for the `afd_configure` command

Option	Description
<code>-f</code>	Force Oracle ASMFd configuration.

The `afd_configure` command must be run noninteractively as a `root` user. Before running the command on a node, the Oracle Clusterware stack must be shut down on that node. After the `afd_configure` command has been run, restart Oracle Clusterware on each node. After running `afd_configure`, the default filtering state is enabled.

Example

The following example configures Oracle ASM Filter Driver on the node on which the command was run. `$ORACLE_HOME` in the example refers to the Oracle Grid Infrastructure home.

Example 10-32 Using the `afd_configure` command

```
# $ORACLE_HOME/bin/asmcmd afd_configure
```

See Also:

- [Administering Oracle ASM Filter Driver](#) (page 20-1) for information about Oracle ASM Filter Driver
- *Oracle Grid Infrastructure Installation and Upgrade Guide* for your operating system for information about installing and configuring Oracle Grid Infrastructure

10.4.2 `afd_deconfigure`

Purpose

Deconfigures Oracle ASM Filter Driver.

Syntax and Description

```
# asmcmd afd_deconfigure [-f]
```

The following table lists the options for the `afd_deconfigure` command.

Table 10-29 Options for the `afd_deconfigure` command

Option	Description
<code>-f</code>	Force Oracle ASMFd deconfiguration.

The `afd_deconfigure` command must be run noninteractively as a `root` user. Before running the command on a node, the Oracle Clusterware stack must be shut down on that node. After the `afd_deconfigure` command has been run, restart Oracle Clusterware on each node.

Example

The following example deconfigures Oracle ASM Filter Driver on the node on which the command was run. `$ORACLE_HOME` in the example refers to the Oracle Grid Infrastructure home.

Example 10-33 Using the `afd_deconfigure` command

```
# $ORACLE_HOME/bin/asmcmd afd_deconfigure
```

See Also:

- [Administering Oracle ASM Filter Driver](#) (page 20-1) for information about Oracle ASM Filter Driver
- *Oracle Grid Infrastructure Installation and Upgrade Guide* for your operating system for information about installing and configuring Oracle Grid Infrastructure

10.4.3 afd_di**Purpose**

Manipulates the data integrity state in the Oracle ASM Filter Driver (Oracle ASMFDF).

Syntax and Description

```
afd_di {-e | -d | -q}
```

The following table lists the options for the `afd_di` command.

Table 10-30 Options for the `afd_di` command

Option	Description
-e	Enables data integrity in Oracle ASMFDF.
-d	Disables data integrity in Oracle ASMFDF.
-q	Queries data integrity in Oracle ASMFDF.

For information about Oracle ASM Filter Driver, refer to "[Administering Oracle ASM Filter Driver](#) (page 20-1)".

Example

The following example shows the use of the `afd_di` command to enable data integrity in the Oracle ASMFDF driver

Example 10-34 Using the ASMCMD `afd_di` command

```
ASMCMD [+] > afd_di -e
```

10.4.4 afd_dsget

Purpose

Retrieves the Oracle ASM Filter Driver discovery diskstring value.

Syntax and Description

```
afd_dsget
```

For information about Oracle ASM Filter Driver, refer to "[Oracle ASM Filter Driver](#) (page 20-1)".

Example

The following example returns the current Oracle ASM Filter Driver discovery diskstring value.

Example 10-35 Using the afd_dsget command

```
ASMCMD [+] > afd_dsget  
AFD discovery string: /dev/rdisk/mydisks/*
```

10.4.5 afd_dsset

Purpose

Sets the Oracle ASM Filter Driver discovery diskstring value.

Syntax and Description

```
afd_dsset afd_diskstring
```

The syntax options for the `afd_dsset` command are described in [Table 10-31](#) (page 10-36).

Table 10-31 Options for the afd_dsset command

Option	Description
<i>afd_diskstring</i>	Specifies the value for the Oracle ASM Filter Driver discovery diskstring.

For information about Oracle ASM Filter Driver, refer to "[Oracle ASM Filter Driver](#) (page 20-1)".

Example

The following example sets the current Oracle ASM Filter Driver discovery diskstring value.

Example 10-36 Using the afd_dsset command

```
ASMCMD [+] > afd_dsset /dev/rdisk/mydisks/*
```

10.4.6 afd_filter

Purpose

Sets the Oracle ASM Filter Driver filtering mode on a given disk path.

Syntax and Description

```
afd_filter {-e | -d } [disk_path]
```

The syntax options for the `afd_filter` command are described in the following table.

Table 10-32 Options for the `afd_filter` command

Option	Description
-e	Enables Oracle ASM Filter Driver mode.
-d	Disables Oracle ASM Filter Driver mode.
<i>disk_path</i>	Specifies the path to the disks.

If the command is executed without specifying a disk path, then filtering is set at the node level.

Example

The following example uses `afd_filter` to enable Oracle ASM Filter Driver filtering on a specified disk path.

Example 10-37 Using the `afd_filter` command

```
ASMCMD [+] > afd_filter -e /dev/sdq
```

10.4.7 `afd_label`

Purpose

Sets an Oracle ASM Filter Driver (Oracle ASMFDD) label to the specified disk.

Syntax and Description

```
afd_label afd_label disk_path [--migrate | --rename] [--init]
```

The syntax options for the `afd_label` command are described in [Table 10-33](#) (page 10-37).

Table 10-33 Options for the `afd_label` command

Option	Description
<i>afd_label</i>	Specifies an Oracle ASM Filter Driver label.
<i>disk_path</i>	Specifies the path to the disks to which the label is applied.
--migrate	Specifies to migrate Oracle ASM Filter Driver disk labels to all disks in the <i>disk_path</i> .
--rename	Specifies to label a disk that was previously labeled.

Table 10-33 (Cont.) Options for the `afd_label` command

Option	Description
<code>--init</code>	Specifies to set an Oracle ASMFD label during the initialize stage. This option is not allowed if Oracle ASMFD is already loaded.

For information about Oracle ASM Filter Driver, refer to "[Administering Oracle ASM Filter Driver](#) (page 20-1)".

Example

The following example sets an Oracle ASM Filter Driver label to a specified disk.

Example 10-38 Using the `afd_label` command

```
ASMCMD [+] > afd_label 'disk0' '/dev/rdsk/mydisks/disk0'
```

10.4.8 `afd_lsdsd`

Purpose

Lists Oracle ASM Filter Driver disks.

Syntax and Description

```
afd_lsdsd [--all]
```

The following table lists the options for the `afd_lsdsd` command.

Table 10-34 Options for the `afd_lsdsd` command

Option	Description
<code>--all</code>	Lists all Oracle ASMFD disks clusterwide.

Example

The following example lists Oracle ASM Filter Driver disks.

Example 10-39 Using the `afd_lsdsd` command

```
ASMCMD [+] > afd_lsdsd
-----
Label           Filtering  Path
-----
DISK0           ENABLED   /dev/sdd
DISK1           DISABLED  /dev/sdm
DISK2           ENABLED   /dev/sdq
```

10.4.9 `afd_lsbl`

Purpose

Scans operating system disks using the specified discovery string and lists the disks that contain Oracle ASM Filter Driver labels.

Syntax and Description

```
afd_ls1bl [disk_string]
```

The following table lists the options for the `afd_ls1bl` command.

Table 10-35 Options for the `afd_ls1bl` command

Option	Description
<code>disk_string</code>	Specifies an optional disk path string for filtering.

For information about Oracle ASM Filter Driver, refer to "[Administering Oracle ASM Filter Driver](#) (page 20-1)".

Example

The following example shows the use of the `afd_ls1bl` command to list disks that contain Oracle ASMFDF labels.

Example 10-40 Using the ASMCMD `afd_ls1bl` command

```
ASMCMD [+]> afd_ls1bl '/dev/sd*'
```

```
Label                Duplicate Path
=====
DISK0                /dev/sdd
DISK1                /dev/sdm
```

10.4.10 `afd_refresh`

Purpose

Refreshes Oracle ASM Filter Driver disks using the AFD discovery string.

Syntax and Description

```
afd_refresh [--all]
```

The following table lists the options for the `afd_refresh` command.

Table 10-36 Options for the `afd_refresh` command

Option	Description
<code>--all</code>	Specifies a clusterwide refresh.

If the Oracle ASM Filter Driver disks are owned by the `root` user, then you must run `afd_refresh --all` as `root` to ensure that accurate disk sizes are displayed in the `$ASM_DISK` view.

For information about Oracle ASM Filter Driver, refer to [Administering Oracle ASM Filter Driver](#) (page 20-1).

Example

The following example uses `afd_refresh` to refresh Oracle ASMFDF disks

Example 10-41 Using the ASMCMD `afd_refresh` command

```
ASMCMD [+]> afd_refresh
```

10.4.11 `afd_scan`

Purpose

Scans for Oracle ASM Filter Driver disks.

Syntax and Description

```
afd_scan [afd_diskstring] [--all]
```

The syntax options for the `afd_scan` command are described in [Table 10-37](#) (page 10-40).

Table 10-37 Options for the `afd_scan` command

Option	Description
<i>afd_diskstring</i>	Specifies the value for the Oracle ASM Filter Driver discovery diskstring.
--all	Specifies a clusterwide scan for Oracle ASMFd disks.

`afd_scan` scans for Oracle ASM Filter Driver disks using the specified discovery string. If the command is executed without specifying a disk string, then `afd_scan` uses the Oracle ASM Filter Driver discovery diskstring value in the `afd.conf` file.

For information about Oracle ASM Filter Driver, refer to "[Oracle ASM Filter Driver](#)" (page 20-1).

Example

The following example scans for Oracle ASM Filter Driver disks with a specified diskstring.

Example 10-42 Using the `afd_scan` command

```
ASMCMD [+] > afd_scan /dev/sd*
```

10.4.12 `afd_state`

Purpose

Returns the state of Oracle ASM Filter Driver.

Syntax and Description

```
afd_state [--all]
```

[Table 10-38](#) (page 10-41) lists the options for the `afd_state` command.

Table 10-38 Options for the `afd_state` command

Option	Description
<code>--all</code>	The <code>--all</code> option runs a clusterwide check.

For information about Oracle ASM Filter Driver, refer to "[Oracle ASM Filter Driver](#) (page 20-1)".

Example

The following example returns the current state of Oracle ASM Filter Driver.

Example 10-43 Using the `afd_state` command

```
ASMCMD [+] > afd_state
ASMCMD-9526: The AFD state is 'LOADED' and filtering is 'DEFAULT' on host 'myhost'
```

10.4.13 `afd_unlabel`

Purpose

Clears a specified Oracle ASM filter driver (Oracle ASMFD) label.

Syntax and Description

```
afd_unlabel afd_label {afd_label | 'disk_path'} [-f] [--init]
```

The syntax options for the `afd_unlabel` command are described in [Table 10-39](#) (page 10-41).

Table 10-39 Options for the `afd_unlabel` command

Option	Description
<code>afd_unlabel</code>	Specifies an existing Oracle ASM Filter Driver label to clear.
<code>'disk_path'</code>	Specifies to clear a label determined by an existing disk path. If a disk needs to be unlabeled during the initialization stage, then it can be done only by using its disk path.
<code>-f</code>	Forces the clearing of the label if the disk is an Oracle ASM member disk and there are no pending opens of the disk.
<code>--init</code>	Specifies to clear an Oracle ASMFD label during the initialize stage. This option is not allowed if Oracle ASMFD is already loaded.

For information about Oracle ASM Filter Driver, refer to "[Administering Oracle ASM Filter Driver](#) (page 20-1)".

Example

The following example clears an existing Oracle ASM Filter Driver label.

Example 10-44 Using the `afd_unlabel` command

```
ASMCMD [+] > afd_unlabel 'disk0'
```

10.5 ASMCMD File Management Commands

This topic provides a summary of the ASMCMD file management commands.

[Table 10-40](#) (page 10-42) lists the ASMCMD file management commands with brief descriptions.

Table 10-40 Summary of ASMCMD file management commands

Command	Description
amdu_extract (page 10-42)	Extracts a file from a disk group.
cd (page 10-43)	Changes the current directory to the specified directory.
cp (page 10-44)	Enables you to copy files between disk groups, and between a disk group and the operating system.
du (page 10-45)	Displays the total disk space occupied by files in the specified Oracle ASM directory and all of its subdirectories, recursively.
find (page 10-46)	Lists the paths of all occurrences of the specified name (with wildcards) under the specified directory.
ls (page 10-47)	Lists the contents of an Oracle ASM directory, the attributes of the specified file, or the names and attributes of all disk groups.
lsof (page 10-49)	Lists the open files.
mkalias (page 10-50)	Creates an alias for system-generated file names.
mkdir (page 10-51)	Creates Oracle ASM directories.
pwd (page 10-51)	Displays the path of the current Oracle ASM directory.
rm (page 10-52)	Deletes the specified Oracle ASM files or directories.
rmbalias (page 10-53)	Deletes the specified alias, retaining the file that the alias points to.

10.5.1 amdu_extract

Purpose

Extracts a file from a disk group.

Syntax and Description

```
amdu_extract diskgroup file_name disk_string
```

Table 10-41 Options for the amdu_extract command

Option	Description
<i>diskgroup</i>	Name of the disk group where the file is located.
<i>file_name</i>	Name of the file to extract. The file name must be an Oracle ASM alias name. The file name can be specified as an absolute or a relative path.

Table 10-41 (Cont.) Options for the *amdu_extract* command

Option	Description
<i>disk_string</i>	Value for the Oracle ASM discovery diskstring.

amdu_extract calls the Oracle ASM Metadata Dump Utility (AMDU) command to extract a file using an Oracle ASM alias name.

Example

The following is an example of the *amdu_extract* command used to extract a file from the data disk group.

Example 10-45 Using the ASMCMD *amdu_extract* command

```
ASMCMD> amdu_extract data data/orcl/my_alias_filename /devices/disk*
```

10.5.2 cd

Purpose

Changes the current directory to the specified directory.

Syntax and Description

```
cd [dir]
```

[Table 10-42](#) (page 10-43) lists the options for the *cd* command.

Table 10-42 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](#)" (page 10-5)".

Examples

The following are examples of the *cd* command changing into various directories.

Example 10-46 Using the ASMCMD *cd* command

```
ASMCMD [+data/hr] > cd +data/orcl
```

```
ASMCMD [+data/orcl] > cd DATAFILE
```

```
ASMCMD [+data/orcl/DATAFILE] >cd ..
```

10.5.3 cp

Purpose

Enables you to copy files between Oracle ASM disk groups and between a disk group and the operating system.

Syntax and Description

```
cp src_file [--target target_type] [--service service_name] [--port port_num]
[connect_str:]tgt_file
```

[Table 10-43](#) (page 10-44) lists the syntax options for the `cp` command.

Table 10-43 Options for the `cp` command

Option	Description
<code>--target <i>target_type</i></code>	Specifies the target type of the instance that ASMCMD must be connected to for the copy operation. Valid options are ASM, IOS, or APX.
<code>--service <i>service_name</i></code>	Specifies the Oracle ASM instance name if not the default +ASM.
<code>--port <i>port_num</i></code>	Specifies the listener port number. The default is 1521.
<code><i>connect_str</i></code>	Specifies the connection string for a remote instance.
<code><i>src_file</i></code>	Name of the source file to copy.
<code><i>tgt_file</i></code>	A user alias for the created target file name or an 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 the OCR and SPFILE file types. To back up, copy, or move an Oracle ASM SPFILE, use the `spbackup`, `spcopy`, or `spmmove` commands. To copy an OCR backup file type, the source must be from a disk group.

The format of copied files is portable between Little-Endian and Big-Endian systems if the files exist in an Oracle ASM disk group. Oracle ASM automatically converts the format when the files are written. For copying a non-Oracle ASM file from or to an Oracle ASM disk group, you can copy the file to a different endian platform and then use one of the commonly used utilities to convert the file.

`connect_str` is not required for a local instance copy, which is the default case. For 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:

user@host.SID

user, *host*, and *SID* are required in the *connect_str* parameter. The default port number is 1521 and can be changed with the `--port` option.

The connection privilege (SYSASM or SYSDBA) is determined by the use of the `--privilege` option when starting ASMCMD.

src_file must be either a fully qualified file name or an Oracle ASM alias.

During an ASMCMD copy process, the Oracle ASM server creates an Oracle Managed File (OMF) file in the following format:

diskgroup/db_unique_name/file_type/file_name.##

where *db_uniq_name* is set to ASM and *#* is a string of digits. In the copy process, ASMCMD `cp` creates the directory structure for the destination and creates a user alias from that location to the actual OMF file that is created.

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. The third example shows how to copy a file in the data disk group to a remote server. You are prompted to enter a password.

Example 10-47 Using the ASMCMD cp command

```
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
```

```
ASMCMD [+] > cp +data/orcl/datafile/EXAMPLE.265.691577295 sys@myserver.+ASM:/scratch/
backup/myexamples.bak
Enter Password:
```

See Also:

- [Specifying the Type of Connection](#) (page 10-9) for more information about connection options with ASMCMD
 - *Oracle Database Net Services Administrator's Guide* for more information about connection strings
-

10.5.4 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 [--suppressheader] [dir]
```

[Table 10-44](#) (page 10-46) lists the syntax options for the `du` command.

Table 10-44 Options for the du command

Option	Description
<i>dir</i>	Name of the directory.
<code>--suppressheader</code>	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](#) (page 10-5)".

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 and each file in the disk group is 2-way mirrored, then `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 the data disk group, including all of the directories under the `orcl` directory.

Example 10-48 Using the ASMCMD du command

```
ASMCMD [+] > du data/orcl
Used_MB      Mirror_used_MB
  1756             3519
```

10.5.5 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 10-45](#) (page 10-46) lists the syntax options for the `find` command.

Table 10-45 Options for the find command

Option	Description
<code>--type type</code>	Type of target 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 used for *pattern* can be a

directory name or a file name, and can include wildcard characters. See "[Wildcard Characters](#) (page 10-5)".

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 file names.

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 5-1](#) (page 5-1). 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 10-49 Using the ASMCMD find command

```
ASMCMD [+] > find +data undo*
+data/ORCL/DATAFILE/UNDOTBS1.258.691577151

ASMCMD [+] > find --type CONTROLFILE +data/orcl *
+data/orcl/CONTROLFILE/Current.260.691577263
```

10.5.6 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 [--suppressheader] [-lsdtLg][--absolutepath][--reverse]
    [--permission][pattern]
```

[Table 10-46](#) (page 10-47) lists the syntax options for the `ls` command.

Table 10-46 Options for the ls command

Option	Description
(none)	Displays only file names 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 <code>SYS</code> column) or user-created (as with an alias, indicated by N under the <code>SYS</code> column). When used in the " <code>ls -l +</code> " 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.

Table 10-46 (Cont.) Options for the ls command

Option	Description
-d	If the value for the <i>pattern</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.
--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>pattern</i> argument is an alias, then ASMCMD displays information about the file that it references. Typically used with another flag, such as the <code>-l</code> flag.
--absolutepath	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.
--suppressheader	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 a pattern.

Command options enable you to modify and customize the output of the command. [Table 10-46](#) (page 10-47) 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.

pattern can be a file name, directory name, or a pattern including wildcard characters. See "[Wildcard Characters](#) (page 10-5)".

- If *pattern* 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 *pattern* is a file name, 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 file name 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 10-50 Using the ASMCMD ls command

```
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    SYSAUX.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
```

10.5.7 lsof

Purpose

Lists the open files of the local clients.

Syntax and Description

```
lsof [--suppressheader] [-G diskgroup | --dbname db | -C instance]
```

[Table 10-47](#) (page 10-49) lists the syntax options for the `lsof` command.

Table 10-47 Options for the `lsof` command

Option	Description
<code>--suppressheader</code>	Suppresses column headings.
<code>-G <i>diskgroup</i></code>	List files only from the specified disk group.
<code>--dbname <i>db</i></code>	List files only from the specified database.
<code>-C <i>instance</i></code>	List files only from the specified instance.

Example

The following are examples of the `ls` 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 10-51 Using the ASMCMD `ls` command

```
ASMCMD [+] > ls -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 [+] > ls -C +ASM
DB_Name Instance_Name Path
asmvol +ASM +data/VOLUME1.271.679226013
asmvol +ASM +data/VOLUME2.272.679227351
```

10.5.8 mkalias

Purpose

Creates an alias for the specified system-generated file name.

Syntax and Description

```
mkalias file alias
```

[Table 10-48](#) (page 10-50) lists the syntax options for the `mkalias` command.

Table 10-48 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 file name `+data/orcl/DATAFILE/SYSAUX.257.721811945`. Following the `mkalias` command, `ls --absolutepath` is run to check the results.

Example 10-52 Using the ASMCMD `mkalias` command

```
ASMCMD [+data/orcl/datafile] > mkalias SYSAUX.257.721811945 sysaux.f

ASMCMD [+data/orcl/datafile] > ls --absolutepath
```



```

none => EXAMPLE.265.721812093
+DATA/ORCL/DATAFILE/sysaux.f => SYSAUX.257.721811945
none => SYSTEM.256.721811945
none => UNDOTBS1.258.721811945
none => USERS.259.721811947
sysaux.f

```

10.5.9 mkdir

Purpose

Creates Oracle ASM directories under the current directory.

Syntax and Description

```
mkdir dir [dir . . .]
```

Table 10-49 (page 10-51) lists the syntax options for the `mkdir` command.

Table 10-49 Options for the `mkdir` command

Option	Description
<code>dir</code>	Directory name to create.

The directory from which `mkdir` is run can be a system-created or user-created directory. 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 10-53 Using the ASMCMD `mkdir` command

```
ASMCMD [+data] > mkdir subdir1 subdir2
```

```

ASMCMD [+data] > ls
ASM/
ORCL/
subdir1/
subdir2/

```

10.5.10 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 10-54 Using the ASMCMD pwd command

```
ASMCMD [+data/orcl/datafile] > pwd
+data/orcl/datafile
```

10.5.11 rm**Purpose**

Deletes the specified Oracle ASM files and directories.

Syntax and Description

```
rm [-f|-r] pattern [pattern...]
```

[Table 10-50](#) (page 10-52) lists the syntax options for the `rm` command.

Table 10-50 Options for the rm command

Option	Description
<code>-r</code>	Recursively deletes files and subdirectories.
<code>-f</code>	Deletes files and subdirectories without prompting for confirmation.
<i>pattern</i>	Name of a file, directory, or wildcard pattern.

If *pattern* is a file or alias, then the `rm` command can delete the file or alias only if it is not currently in use. If *pattern* 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 *pattern* 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.

pattern can contain wildcard characters. See "[Wildcard Characters](#) (page 10-5)".

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 to delete an alias that matches a wildcard pattern, either the alias or the system-generated file that has an 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 10-55 Using the ASMCMD rm command

```
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
```

10.5.12 rmalias

Purpose

Removes the specified aliases, retaining the files that the aliases reference.

Syntax and Description

```
rmalias [-r] alias [alias...]
```

[Table 10-51](#) (page 10-53) lists the syntax options for the rmalias command.

Table 10-51 Options for the rmalias command

Option	Description
-r	Recursively removes aliases.
alias	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 after 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 10-56 Using the ASMCMD rmalias command

```
ASMCMD [+data/orcl/datafile] > rmalias sysaux.f
```

10.6 ASMCMD Disk Group Management Commands

This topic provides a summary of the ASMCMD disk group management commands.

[Table 10-52](#) (page 10-54) lists the ASMCMD disk group management commands with brief descriptions.

Table 10-52 Summary of ASMCMD Disk group management commands

Command	Description
chdg (page 10-54)	Changes a disk group (add, drop, or rebalance).
chkdg (page 10-57)	Checks or repairs a disk group.
dropdg (page 10-57)	Drops a disk group.
iostat (page 10-58)	Displays I/O statistics for disks.
lsattr (page 10-60)	Lists the attributes of a disk group.
lsdg (page 10-61)	Lists disk groups and their information.
lsdisk (page 10-63)	Lists disks Oracle ASM disks.
lsod (page 10-65)	Lists open devices.
md_backup (page 10-66)	Creates a backup of the metadata of mounted disk groups.
md_restore (page 10-67)	Restores disk groups from a backup of the metadata.
mkdg (page 10-69)	Creates a disk group.
mount (page 10-71)	Mounts a disk group.
offline (page 10-72)	Offlines a disk or a failure group.
online (page 10-73)	Onlines a disk or a failure group.
rebal (page 10-74)	Rebalances a disk group.
remap (page 10-75)	Relocates data in a range of physical blocks on a disk.
setattr (page 10-75)	Sets attributes in a disk group.
stamp (page 10-76)	Stamps the disk, site, and failure group labels in disk headers.
stamplist (page 10-77)	Displays the disk, site, and failure group labels in disk headers.
umount (page 10-78)	Dismounts a disk group.

10.6.1 chdg

Purpose

Changes a disk group (adds disks, drops disks, resizes disks, or rebalances a disk group) based on an XML configuration file.

Syntax and Description

```
chdg { config_file.xml | 'contents_of_xml_file' }
```

[Table 10-53](#) (page 10-55) lists the syntax options for the `chdg` command.

Table 10-53 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 10-57 (page 10-55) and Example 10-58 (page 10-56).
<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, replacing, or deleting disks from an existing disk group, and the setting rebalance or replace power level. The power level can be set to the same values as the ASM_POWER_LIMIT initialization parameter.

When adding disks to a disk group, the diskstring must be specified in a format similar to the ASM_DISKSTRING initialization parameter.

The failure groups are optional parameters. The default causes every disk to belong to a its own failure group.

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 set of disks that belong to a failure group can be specified by the failure group name.

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.

[Example 10-57](#) (page 10-55) shows the basic structure and the valid tags with their respective attributes for the chdg XML configuration file.

Example 10-57 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 or replace

<add>  items to add are placed here
</add>

<replace> items to replace are placed here
</replace>

<drop> items to drop are placed here
</drop>

<fg>  failure group
      name      failure group name
</fg>

<dsk> disk
      name      disk name
      string    disk path
      size      size of the disk to add
      force     true specifies to use the force option
</dsk>

</chdg>
```

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 10-58 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 10-59 Using the ASMCMD chdg command

```
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>'
```

See Also:

- [ASM_POWER_LIMIT](#) (page 3-11) for information about the `ASM_POWER_LIMIT` initialization parameter
 - [ASM_DISKSTRING](#) (page 3-10) for information about the `ASM_DISKSTRING` initialization parameter
 - [Oracle ASM Failure Groups](#) (page 4-40) for information about failure groups
 - [Dropping Disks from Disk Groups](#) (page 4-25) for information about dropping disks
 - [Resizing Disks in Disk Groups](#) (page 4-27) for information about resizing disks,
 - [mkdg](#) (page 10-69) for information about creating a disk group with ASMCMD `mkdg`
 - [Altering Disk Groups](#) (page 4-19) for information about altering disk groups
-

10.6.2 chkdg

Purpose

Checks or repairs the metadata of a disk group.

Syntax and Description

```
chkdg [--repair] diskgroup
```

[Table 10-54](#) (page 10-57) lists the syntax options for the `chkdg` command.

Table 10-54 Options for the `chkdg` command

Option	Description
<code>--repair</code>	Repairs the disk group.
<code>diskgroup</code>	Name of disk group to check or repair.

`chkdg` checks the metadata of a disk group for errors and optionally repairs the errors.

Example

The following is an example of the `chkdg` command used to check and repair the `data` disk group.

Example 10-60 Using the ASMCMD `chkdg` command

```
ASMCMD [+] > chkdg --repair data
```

10.6.3 dropdg

Purpose

Drops a disk group.

Syntax and Description

```
dropdg [-r [-f]] diskgroup
```

[Table 10-55](#) (page 10-57) lists the syntax options for the `dropdg` command.

Table 10-55 Options for the `dropdg` command

Option	Description
<code>-f</code>	Force the operation. Only applicable if the disk group cannot be mounted.
<code>-r</code>	Recursive, include contents.
<code>diskgroup</code>	Name of disk group to drop.

`dropdg` drops an existing disk group. The disk group should not be mounted on multiple nodes.

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 10-61 Using the ASMCMD dropdg command

```
ASMCMD [+] > dropdg -r -f data
```

```
ASMCMD [+] > dropdg -r fra
```

10.6.4 iostat

Purpose

Displays I/O statistics for Oracle ASM disks in mounted disk groups.

Syntax and Description

```
iostat [--suppressheader] [-et] [--io] [--region]
      [-G diskgroup] [interval]
```

`iostat` lists disk group statistics using the `V$ASM_DISK_IOSTAT` view.

[Table 10-56](#) (page 10-58) lists the syntax options for the `iostat` command.

Table 10-56 Options for the iostat command

Option	Description
<code>-e</code>	Displays error statistics (Read_Err, Write_Err).
<code>-G <i>diskgroup</i></code>	Displays statistics for the disk group name.
<code>--suppressheader</code>	Suppresses column headings.
<code>--io</code>	Displays information in number of I/Os, instead of bytes.
<code>-t</code>	Displays time statistics (Read_Time, Write_Time).
<code>--region</code>	Displays information for cold and hot disk regions (Cold_Reads, Cold_Writes, Hot_Reads, Hot_Writes).
<code><i>interval</i></code>	Refreshes the statistics display based on the interval value (seconds). Use CTRL+C to stop the interval display.

[Table 10-57](#) (page 10-58) 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 10-57 Attribute descriptions for iostat command output

Attribute Name	Description
Group_Name	Name of the disk group.
Dsk_Name	Name of the disk.
Reads	Number of bytes read from the disk. If the <code>--io</code> option is entered, then the value is displayed as number of I/Os.

Table 10-57 (Cont.) Attribute descriptions for iostat command output

Attribute Name	Description
Writes	Number of bytes written to the disk. If the <code>--io</code> option is entered, then the value is displayed as number of I/Os.
Cold_Reads	Number of bytes read from the cold disk region. If the <code>--io</code> option is entered, then the value is displayed as number of I/Os.
Cold_Writes	Number of bytes written from the cold disk region. If the <code>--io</code> option is entered, then the value is displayed as number of I/Os.
Hot_Reads	Number of bytes read from the hot disk region. If the <code>--io</code> option is entered, then the value is displayed as number of I/Os.
Hot_Writes	Number of bytes written to the hot disk region. If the <code>--io</code> option is entered, then the value is displayed as number of I/Os.
Read_Err	Number of failed I/O read requests for the disk.
Write_Err	Number of failed I/O write requests for the disk.
Read_Time	I/O time (in seconds) for read requests for the disk if the <code>TIMED_STATISTICS</code> initialization parameter is set to <code>TRUE</code> (0 if set to <code>FALSE</code>).
Write_Time	I/O time (in seconds) for write requests for the disk if the <code>TIMED_STATISTICS</code> initialization parameter is set to <code>TRUE</code> (0 if set to <code>FALSE</code>).

If a refresh interval is not specified, the number displayed represents the total number of bytes or I/Os. If a refresh interval is specified, then the value displayed (bytes or I/Os) is the difference between the previous and current values, not the total value.

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 10-62 Using the ASMCMD iostat command

```

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
    
```

10.6.5 lsattr

Purpose

Lists the attributes of a disk group.

Syntax and Description

```
lsattr [--suppressheader][-G diskgroup ] [-lm] [pattern]
```

[Table 10-58](#) (page 10-60) lists the syntax options for the `lsattr` command.

Table 10-58 Options for the `lsattr` command

Option	Description
<code>-G <i>diskgroup</i></code>	Disk group name.
<code>--suppressheader</code>	Suppresses column headings.
<code>-l</code>	Display names with values.
<code>-m</code>	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 "[Managing Disk Group Attributes](#) (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](#) (page 10-88)".

To set disk group attributes, see "[setattr](#) (page 10-75)".

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. Note the use of both the `%` and `*` wildcard characters on Linux.

Example 10-63 Using the ASMCMD lsattr command

```

ASMCMD [+] > lsattr -lm -G data
Group_Name Name Value RO Sys
DATA access_control.enabled FALSE N Y
DATA access_control.umask 066 N Y
DATA au_size 1048576 Y Y
DATA cell.smart_scan_capable FALSE N N
DATA cell.sparse_dg allnonsparse N N
DATA compatible.advm 12.1.0.2.0 N Y
DATA compatible.asm 12.1.0.2.0 N Y
DATA compatible.rdbms 12.1.0.0.0 N Y
DATA content.check FALSE N Y
DATA content.type data N Y
DATA disk_repair_time 3.6h N Y
DATA failgroup_repair_time 24.0h N Y
DATA idp.boundary auto N Y
DATA idp.type dynamic N Y
DATA phys_meta_replicated true Y Y
DATA sector_size 512 Y Y
DATA thin_provisioned FALSE N Y

ASMCMD [+] > lsattr -G fra -l %compat*
Name Value
compatible.asm 12.1.0.0.0
compatible.rdbms 11.2.0.3.0
    
```

10.6.6 lsdg

Purpose

Lists mounted 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 [--suppressheader] [-g] [--discovery] [pattern]
```

[Table 10-59](#) (page 10-61) lists the options for the `lsdg` command.

Table 10-59 Options for the lsdg command

Option	Description
(none)	Displays the disk group attributes listed in Table 10-60 (page 10-62).
<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>lsdg</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.

Table 10-59 (Cont.) Options for the *lsdg* command

Option	Description
<code>--suppressheader</code>	Suppresses column headings.
<code>pattern</code>	Returns only information about the specified disk group or disk groups that match the supplied pattern. See " Wildcard Characters (page 10-5)".

Table 10-60 (page 10-62) 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 10-60 Attribute descriptions for *lsdg* command output

Attribute Name	Description
State	State of the disk group. Values include <code>BROKEN</code> , <code>CONNECTED</code> , <code>DISMOUNTED</code> , <code>MOUNTED</code> , <code>QUIESCING</code> , and <code>UNKNOWN</code> .
Type	Disk group redundancy (<code>NORMAL</code> , <code>HIGH</code> , <code>EXTERNAL</code>).
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 <code>V\$ASM_DISKGROUP</code> 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 <code>REQUIRED_MIRROR_FREE_MB</code> column from the <code>V\$ASM_DISKGROUP</code> view.
Usable_file_MB	Amount of free space, adjusted for mirroring, that is available for new files. From the <code>V\$ASM_DISKGROUP</code> 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 files (Y or N).
Name	Disk group name.

Example

The following example lists the attributes of the data disk group.

Example 10-64 Using the ASMCMD lsdg command

```
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
```

10.6.7 lsdsk

Purpose

Lists Oracle ASM disks.

Syntax and Description

```
lsdsk [--suppressheader] [-kptgMI] [-G diskgroup ]
      [ --member|--candidate]
      [--discovery][--statistics][pattern]
```

Table 10-61 (page 10-63) lists the options for the lsdsk command.

Table 10-61 Options for the lsdsk command

Option	Description
(none)	Displays the PATH column of the V\$ASM_DISK_STAT view.
-k	Displays the TOTAL_MB, FREE_MB, OS_MB, NAME, FAILGROUP, LIBRARY, LABEL, UDID, PRODUCT, REDUNDANCY, and PATH columns of the V\$ASM_DISK view.
--statistics	Displays the READS, WRITES, READ_ERRS, WRITE_ERRS, READ_TIME, WRITE_TIME, BYTES_READ, BYTES_WRITTEN, and the PATH columns of the V\$ASM_DISK view.
-p	Displays the GROUP_NUMBER, DISK_NUMBER, INCARNATION, MOUNT_STATUS, HEADER_STATUS, MODE_STATUS, STATE, and the PATH columns of the V\$ASM_DISK view.
-t	Displays the CREATE_DATE, MOUNT_DATE, REPAIR_TIMER, and the PATH columns of the V\$ASM_DISK view.
-g	Selects from GV\$ASM_DISK_STAT, or from GV\$ASM_DISK if the --discovery flag is also specified. GV\$ASM_DISK.INST_ID is included in the output.
--discovery	Selects from V\$ASM_DISK, or from GV\$ASM_DISK if the -g 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 lsdsk is running in non-connected mode.
--suppressheader	Suppresses column headings.

Table 10-61 (Cont.) Options for the `lsdisk` command

Option	Description
<code>-I</code>	Scans disk headers for information rather than extracting the information from an Oracle ASM instance. This option forces non-connected mode.
<code>-G</code>	Restricts results to only those disks that belong to the group specified by <i>diskgroup</i> .
<code>-M</code>	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.
<code>--candidate</code>	Restricts results to only disks having membership status equal to CANDIDATE.
<code>--member</code>	Restricts results to only disks having membership status equal to MEMBER.
<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](#) (page 10-5)".

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 10-65 Using the ASMCMD lsdisk command

```

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
...

```

10.6.8 lsod

Purpose

Lists the open Oracle ASM disks.

Syntax and Description

```
lsod [--suppressheader] [-G diskgroup] [--process process] [pattern]
```

Table 10-62 (page 10-65) lists the syntax options for the lsod command.

Table 10-62 Options for the lsod command

Option	Description
--suppressheader	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 10-66 Using the ASMCMD `lsod` command

```
ASMCMD [+] > lsod -G data --process *LGWR*
Instance Process                OSPID Path
1         oracle@myserver02 (LGWR) 26593 /devices/diska1
1         oracle@myserver02 (LGWR) 26593 /devices/diska2
1         oracle@myserver02 (LGWR) 26593 /devices/diska3
1         oracle@myserver02 (LGWR) 26593 /devices/diskb1
1         oracle@myserver02 (LGWR) 26593 /devices/diskb2
1         oracle@myserver02 (LGWR) 26593 /devices/diskb3
1         oracle@myserver02 (LGWR) 26593 /devices/diskd1

ASMCMD [+] > lsod --process *LGWR* *diska*
Instance Process                OSPID Path
1         oracle@myserver02 (LGWR) 26593 /devices/diska1
1         oracle@myserver02 (LGWR) 26593 /devices/diska2
1         oracle@myserver02 (LGWR) 26593 /devices/diska3
```

For another example of the `lsod` command, see [Example 10-2](#) (page 10-10).

10.6.9 md_backup

Purpose

The `md_backup` command creates a backup file containing metadata for one or more disk groups.

Syntax and Description

```
md_backup backup_file
        [-G diskgroup [,diskgroup,...]] --acfs_sec_encr
```

[Table 10-63](#) (page 10-66) describes the options for the `md_backup` command.

Table 10-63 Options for the `md_backup` command

Option	Description
<i>backup_file</i>	Specifies the backup file in which you want to store the metadata.
-G <i>diskgroup</i>	Specifies the disk group name of the disk group that must be backed up
--acfs_sec_encr	Specifies that security, encryption, or audit metadata information from Oracle ACFS is backed up.

By default all the mounted disk groups are included in the backup file, which is saved in the current working directory if a path is not specified with the file name.

The `md_backup` command backs up Oracle Automatic Storage Management Cluster File System (Oracle ACFS) metadata information for the contained Oracle ASM Dynamic Volume Manager(Oracle ADVM) volumes, including stripe size,

redundancy, and other items. This command also backs up Oracle ACFS file system configuration metadata, such as compression and resize attributes, block size, and created snapshots. Additional metadata and configuration attributes may be backed up; however, Oracle ACFS file system data and Oracle ACFS CRS Resource information is not backed up.

Example

The first example shows the use of the backup command when run without the disk group option. This example backs up all the mounted disk groups and creates the backup image in the `/scratch/backup/allldgs20100422` file. The second example creates a backup of the data disk group. The metadata backup that this example creates is saved in the `/scratch/backup/data20100422` file.

Example 10-67 Using the ASMCMD `md_backup` command

```
ASMCMD [+] > md_backup /scratch/backup/allldgs20100422
Disk group metadata to be backed up: DATA
Disk group metadata to be backed up: FRA
Current alias directory path: ORCL/ONLINELOG
Current alias directory path: ORCL/PARAMETERFILE
Current alias directory path: ORCL
Current alias directory path: ASM
Current alias directory path: ORCL/DATAFILE
Current alias directory path: ORCL/CONTROLFILE
Current alias directory path: ASM/ASMPARAMETERFILE
Current alias directory path: ORCL/TEMPFILE
Current alias directory path: ORCL/ARCHIVELOG/2010_04_20
Current alias directory path: ORCL
Current alias directory path: ORCL/BACKUPSET/2010_04_21
Current alias directory path: ORCL/ARCHIVELOG/2010_04_19
Current alias directory path: ORCL/BACKUPSET/2010_04_22
Current alias directory path: ORCL/ONLINELOG
Current alias directory path: ORCL/BACKUPSET/2010_04_20
Current alias directory path: ORCL/ARCHIVELOG
Current alias directory path: ORCL/BACKUPSET
Current alias directory path: ORCL/ARCHIVELOG/2010_04_22
Current alias directory path: ORCL/DATAFILE
Current alias directory path: ORCL/CONTROLFILE
Current alias directory path: ORCL/ARCHIVELOG/2010_04_21

ASMCMD [+] > md_backup /scratch/backup/data20100422 -G data
Disk group metadata to be backed up: DATA
Current alias directory path: ORCL/ONLINELOG
Current alias directory path: ASM
Current alias directory path: ORCL/CONTROLFILE
Current alias directory path: ASM/ASMPARAMETERFILE
Current alias directory path: ORCL/PARAMETERFILE
Current alias directory path: ORCL
Current alias directory path: ORCL/DATAFILE
Current alias directory path: ORCL/TEMPFILE
```

10.6.10 `md_restore`

Purpose

The `md_restore` command restores disk groups from a metadata backup file.

Syntax and Description

```
md_restore backup_file [--silent]
  [--full|--nodg|--newdg -o 'old_diskgroup:new_diskgroup [,...]]
  [--acfs_sec_encr user:group |
  --acfs_sec_encr user:group --acfs_audit user:mgr_group:auditor_group ]
  [-S sql_script_file] [-G 'diskgroup [,diskgroup...]]'
```

Table 10-64 (page 10-68) describes the options for the `md_restore` command.

Table 10-64 Options for the `md_restore` command

Option	Description
<code>backup_file</code>	Reads the metadata information from <code>backup_file</code> .
<code>--silent</code>	Ignore errors. Typically, if <code>md_restore</code> encounters an error, it stops. Specifying this flag ignores any errors.
<code>--full</code>	Specifies to create a disk group and restore metadata.
<code>--nodg</code>	Specifies to restore metadata only.
<code>--newdg -o old_diskgroup:new_diskgro up]</code>	Specifies to create a disk group with a different name when restoring metadata. The <code>-o</code> option is required with <code>--newdg</code> .
<code>--acfs_sec_encr user:group</code>	Restores the metadata of security or encryption or both. The colon separated user and group is used for security metadata restoration. During metadata restore, security or encryption or both must be initialized. If security or encryption or both are not pre-initialized, the colon separated values passed to the <code>--acfs_sec_encr</code> option are used to initialize security. Encryption is initialized with the SSO wallet.
<code>--acfs_audit user:mgr_group:auditor_gr oup</code>	Restores audit metadata information on an Oracle ACFS file system. The colon separated values are used to initialize and execute audit commands. During metadata restore, audit must be initialized. If audit is not pre-initialized, the audit layer is initialized using the colon separated values given to the <code>--acfs_audit</code> option. The <code>--acfs_audit</code> option requires the <code>--acfs_sec_encr</code> option.
<code>-S sql_script_file</code>	Write SQL commands to the specified SQL script file instead of executing the commands.
<code>-G diskgroup</code>	Select the disk groups to be restored. If no disk groups are defined, then all disk groups are restored.

Note the following when restoring Oracle ACFS information:

- When restoring the metadata of a file system, the user id is set to the user id that owned the file system when the metadata backup was created. This user id should exist on the restore cluster.
- `md_restore` restores Oracle ACFS metadata information. To run the metadata restore for Oracle ACFS file systems, a script is created. This script must be run as `root` to format file systems. Metadata for the contained Oracle ADVM volumes, such as stripe size and redundancy, is restored. Oracle ACFS file system

configuration metadata, such as compression and resize attributes, block size, and created snapshots, is also restored. Oracle ACFS file System data and ACFS CRS Resource information is not restored.

- If using the `--silent` option while specifying a single disk group, and there is a dependency due to Oracle ADVM metadata or Oracle ACFS accelerators on another disk group, that dependency is dropped, and you must manually recreate that dependency later.
- When restoring Oracle ACFS encryption or security, file system extended attributes (xattrs) must be restored with the user data. It is very important to document that xattrs need to be stored and restored with the file system application data. Without the xattrs, there is no security or encryption. xattrs describe which security realms the file is a part of. With no xattr, the file ends up in no realm. The converse is that a file with xattrs, but no security initialized, ends up in no realm.
- During metadata restore, a file with a list of previously configured snapshots is created. Snapshot data is not restored, as the point in time sparse snapshot views cannot be recreated.
- If Oracle ACFS security is in use, and the metadata security information for the file system has not been backed up, this information may be copied by the security administrator to the metadata backup file location, and then restored with the rest of the disk group information.
- If Oracle ACFS security and encryption is being restored, the user and password may be pre-initialized. If it is not, the username provided is used to initialize Oracle ACFS security and encryption, with a default password which is displayed.

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 named `data2`. The fourth example restores from the backup file after applying the overrides defined in the `override.sql` script file.

Example 10-68 Using the ASMCMD `md_restore` command

```
ASMCMD [+] > md_restore --full -G data --silent /scratch/backup/alldgs20100422

ASMCMD [+] > md_restore --nodg -G data --silent /scratch/backup/alldgs20100422

ASMCMD [+] > md_restore --newdg -o 'data:data2' --silent /scratch/backup/data20100422

ASMCMD [+] > md_restore -S override.sql --silent /scratch/backup/data20100422
```

10.6.11 mkdg

Purpose

Creates a disk group based on an XML configuration file.

Syntax and Description

```
mkdg { config_file.xml | 'contents_of_xml_file' }
```

Table 10-65 (page 10-70) lists the syntax options for the `mkdg` command.

Table 10-65 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.
<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 for a disk group, each disk in the disk group belongs to its own failure group.

The `mkdg` command only mounts a disk group on the local node.

It is possible to set some 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.

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.

Example 10-69 (page 10-70) shows the basic structure and the valid tags with their respective attributes for the `mkdg` XML configuration file.

Example 10-69 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>

<dsk> disk
      name          disk name
      string        disk path
      size          size of the disk to add
      force         true specifies to use the force option
</dsk>

<a>   attribute
      name          attribute name
      value         attribute value
</a>

</dg>
```

Example

The following is an example of an XML configuration file for `mkdg`. 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 10-70 mkdg 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 `mkdg` command. The first example runs `mkdg` with an XML configuration file in the directory where ASMCMD was started. The second example runs `mkdg` using information on the command line.

Example 10-71 Using the ASMCMD mkdg command

```
ASMCMD [+] > mkdg data_config.xml
```

```
ASMCMD [+] > mkdg '<dg name="data"><dsk string="/dev/disk*"/></dg>'
```

See Also:

- [Example 10-69](#) (page 10-70) and [Example 10-70](#) (page 10-71) for examples of the valid tags and XML configuration file
- [Managing Disk Group Attributes](#) (page 4-2) for more information about disk groups attributes
- [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility attributes
- [chdg](#) (page 10-54) for information about altering a disk group with ASMCMD `chdg`
- [Creating Disk Groups](#) (page 4-9) for information about creating a disk group

10.6.12 mount

Purpose

Mounts a disk group.

Syntax and Description

```
mount [--restrict] { [-a] | [-f] diskgroup[ diskgroup ...] }
```

[Table 10-66](#) (page 10-72) lists the syntax options for the `mount` command.

Table 10-66 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](#) (page 4-67)".

Example

The following are examples of the mount command showing the use of the force, restrict, and all options.

Example 10-72 Using the ASMCMD mount command

```
ASMCMD [+] > mount -f data
```

```
ASMCMD [+] > mount --restrict data
```

```
ASMCMD [+] > mount -a
```

10.6.13 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 10-67](#) (page 10-72) lists the syntax options for the `offline` command.

Table 10-67 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>m</i> or <i>nh</i> , where <i>m</i> specifies minutes and <i>h</i> specifies hours. For example, 120m or 2h. 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 10-73 Using the ASMCMD offline command

```
ASMCMD [+] > offline -G data -F failgroup1
ASMCMD [+] > offline -G data -D data_0001 -t 1.5h
```

10.6.14 online

Purpose

Online all disks, a single disk, or a failure group that belongs to a disk group.

Syntax and Description

```
online -G diskgroup { -a | -F failgroup |-D disk }
        [--power n] [-w]
```

Table 10-68 (page 10-73) lists the syntax options for the `online` command.

Table 10-68 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.
--power <i>n</i>	Specifies the power option (1 to 1024) which determines the level of resources for the disk resync and disk group rebalance. Larger values result in faster disk resync and disk group rebalance. The default is 1.
-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 10-74 Using the ASMCMD online command

```
ASMCMD [+] > online -G data -F failgroup1 -w
ASMCMD [+] > online -G data -D data_0001
```

10.6.15 rebal

Purpose

Rebalances a disk group.

Syntax and Description

```
rebal [--power power] [-w] diskgroup
```

[Table 10-69](#) (page 10-74) lists the syntax options for the `rebal` command.

Table 10-69 Options for the rebal command

Option	Description
<i>diskgroup</i>	Disk group name.
<code>--power power</code>	Power setting.
<code>-w</code>	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 to the same values as the `ASM_POWER_LIMIT` initialization parameter. 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.

You can determine if a rebalance operation is occurring with the ASMCMD `lsop` command.

Example

The following is an example of the `rebal` command that rebalances the `fra` disk group with a power level set to 4.

Example 10-75 Using the ASMCMD rebal command

```
ASMCMD [+] > rebal --power 4 fra

ASMCMD [+] > lsop
Group_Name Dsk_Num State Power
FRA        REBAL  RUN    4
```


See Also:

- [ASM_POWER_LIMIT](#) (page 3-11) and [Tuning Rebalance Operations](#) (page 4-31) for information about the power level
- [lsop](#) (page 10-15) for information about the ASMCMD `lsop` command
- [Manually Rebalancing Disk Groups](#) (page 4-28) for more information about rebalancing a disk group

10.6.16 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 10-70](#) (page 10-75) lists the syntax options for the `remap` command.

Table 10-70 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 <code>NAME</code> column in the <code>V\$ASM_DISK</code> 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 10-76 Using the ASMCMD remap command

```
ASMCMD [+] > remap DATA DATA_0001 5000-5999
```

```
ASMCMD [+] > remap FRA FRA_0002 6230-6339
```

10.6.17 setattr

Purpose

Sets the attributes for an Oracle ASM disk group.

Syntax and Description

```
setattr -G diskgroup attribute_name attribute_value
```

[Table 10-71](#) (page 10-76) lists the syntax options for the `setattr` command.

Table 10-71 Options for the `setattr` command

Option	Description
<code>-G diskgroup</code>	Disk group name.
<code>attribute_name</code>	Name of the attribute.
<code>attribute_value</code>	Value of the attribute.

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.

For information about disk group attributes, see "[Managing Disk Group Attributes](#) (page 4-2)".

Example

The following are examples of the `setattr` command. The first example sets the disk group attribute `COMPATIBLE.ASM` to `12.1.0.2` for the data disk group. The second example sets the disk group attribute `COMPATIBLE.RDBMS` to `12.1` for the data disk group.

Example 10-77 Using the ASMCMD `setattr` command

```
ASMCMD [+] > setattr -G data compatible.asm 12.1.0.2
```

```
ASMCMD [+] > setattr -G data compatible.rdbms 12.1
```

10.6.18 stamp

Purpose

Stamps the disk, site, and failure group labels in disk headers.

Syntax and Description

```
stamp --dscstr disk_discovery_string
      [--site site_name --failgroup failgroup_name]
      [--disk disk_label] [-f]
```

The following table describes the options for the `stamp` command.

Table 10-72 Options for the `stamp` command

Option	Description
<code>--dscstr disk_discovery_string</code>	Specifies the disk discovery string that identifies the disks on which to stamp labels.
<code>--site site_name</code>	Specifies the site label name.

Table 10-72 (Cont.) Options for the stamp command

Option	Description
--failgroup <i>failgroup_name</i>	Specifies the failure group label. Required when the site label is specified.
--disk <i>disk_label</i>	Specifies the disk label.
-f	Specifies to force the action.

ASMCMD `stamp` enables the users to stamp the disk, site, and fail group labels on disk headers. Using `stamp` is more efficient than individually specifying label information during installation and creation of disk groups.

Site label and fail groups labels are used for the site names and fail group names when creating and altering a disk group. Disk labels are only intended for use with Oracle ASM filter driver (AFD) and Oracle ASM libraries.

At least one label specification (`--site` with `--failgroup` or `--disk`) must be provided.

Examples

The following example shows the use of the ASMCMD `stamp` command on the disks identified by the `/dev/sd*` disk discovery string.

Example 10-78 Using ASMCMD stamp

```
ASMCMD [+] > stamp --dscstr "/dev/sd*" --site "SALT LAKE-CENTER" --failgroup "fg1" --disk "DISK1"
```

10.6.19 stamp list

Purpose

Displays the disk, site and failure group labels in disk headers.

Syntax and Description

```
stamp --dscstr disk_discovery_string {[--site] [--failgroup] [--disk]}
```

The following table describes the options for the `stamp list` command.

Table 10-73 Options for the stamp list command

Option	Description
-- <i>dscstrdisk_discovery_string</i>	Specifies the disk discovery string that identifies the disks for which to display labels.
--site	Displays the site label name.
--failgroup	Displays the failure group label.
--disk	Displays the disk label.

You must specify at least one of the `--site`, `--failgroup`, or `--disk` options.

Examples

The following example shows the use of the ASMCMD `stamp` command to display site, failure group, and disk labels of the disks identified by the `/dev/sd*` disk discovery string.

Example 10-79 Using ASMCMD stamp

```
ASMCMD [+] > stamp --dscstr "/dev/sd*" --site --failgroup --disk
```

10.6.20 umount

Purpose

Dismounts a disk group.

Syntax and Description

```
umount { -a | [-f] diskgroup }
```

[Table 10-74](#) (page 10-78) lists the syntax options for the `umount` command.

Table 10-74 Options for the umount command

Option	Description
<i>diskgroup</i>	Name of the disk group.
-a	Dismounts all mounted disk groups. These disk groups are listed in the output of the <code>V\$ASM_DISKGROUP</code> view.
-f	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 10-80 Using the ASMCMD umount command

```
ASMCMD [+] > umount -a
```

```
ASMCMD [+] > umount -f data
```

10.7 ASMCMD File Group Management Commands

This topic provides a summary of the ASMCMD commands that enable you to manage Oracle ASM file groups and quota groups.

The following table lists the ASMCMD file group and quota group commands with brief descriptions.

Table 10-75 Summary of ASMCMD file group commands

Command	Description
chfg (page 10-79)	Modifies a file group in a disk group based on an XML configuration file or script.
chqg (page 10-81)	Modifies a quota group in the disk group.
lsfg (page 10-81)	List file groups in a disk group.
lsqg (page 10-82)	Lists the quota groups in a disk group.
mkfg (page 10-83)	Creates a file group in a disk group based on an XML configuration file or script.
mkqg (page 10-85)	Adds a quota group to the disk group.
mvfg (page 10-85)	Moves a file group in a disk group to the specified quota group.
rmfg (page 10-86)	Removes an existing file group from a disk group.
rmqg (page 10-86)	Drops a quota group from a disk group.

See Also:

- [Managing Oracle ASM Flex Disk Groups](#) (page 4-72) for more information about Oracle file groups
 - *Oracle Database Reference* for information about the V\$ASM views that contain file group and quota group information.
 - *Oracle Database SQL Language Reference* for information about administering Oracle ASM file groups and quota groups with SQL statements
-
-

10.7.1 chfg

Purpose

Changes the attributes of a file group in a disk group based on an XML configuration file or script.

Syntax and Description

```
chfg { config_file.xml | 'contents_of_xml_script' }
```

The following table lists the options for the `chfg` command.

Table 10-76 Options for the chfg command

Option	Description
<code>config_file.xml</code>	Name of the XML file that contains the changes for the file group. <code>chfg</code> searches for the XML file in the current working directory of ASMCMD unless a path is specified
<code>contents_of_xml_script.xml</code>	The contents of an XML script enclosed in single quotes.

The root element of the XML configuration file or script is the `<file_group>` tag and has these attributes:

- `name`: file group name
- `dg`: disk group name to which the file group belongs

You can set properties using these tag:

- `name`: property name
- `value`: property value
- `file_type`: property file type (optional)
- `template`: Oracle ASM template name (optional)

The following is an example of an XML configuration file for `chfg`. The configuration file alters a file group named `filegroup1` on disk group `dg_data1`. The file group is associated with quota group `quotagroup1`. The files in this file group have the redundancy set to mirror, except for the control file, which has a high redundancy. All other attributes are not changed.

```
<filegroup name="filegroup1" dg="dg_data1">
  <p name="redundancy" value="mirror"/>
  <p name="redundancy" value="high" file_type="controlfile"/>
  <p name="quota_group" value="quotagroup1"/>
</filegroup>
```

For information about Oracle ASM File Groups, refer to "[Managing Oracle ASM Flex Disk Groups](#) (page 4-72)".

Example

The following examples show the use of the `chfg` command using an XML configuration file and an XML script.

Example 10-81 Using the chfg command

```
ASMCMD [+] > chfg config_file.xml
```

```
ASMCMD [+] > chfg '<filegroup name="filegroup1" dg="dg_data1">
  <p name="redundancy" value="high"/>
</filegroup>'
```

10.7.2 chqg

Purpose

Modifies a quota group in the disk group.

Syntax and Description

```
chqg -G disk_group quota_group property value
```

The following table lists the options for the chqg command.

Table 10-77 Options for the chqg command

Option	Description
<i>-G disk_group</i>	Specifies the disk group.
<i>quota_group</i>	Specifies the quota group.
<i>property</i>	Specifies the name of the property.
<i>value</i>	Specifies the value of the property.

The chqg command modifies the specified property in the specified quota group of the specified disk group. For information about Oracle ASM File and Quota Groups, refer to "[Managing Oracle ASM Flex Disk Groups](#) (page 4-72)".

The following are examples of the chqg command. The first example modifies the quota property of the quota group *quota_grp1* in the *data* disk group. The second example modifies the quota property of the quota group *quota_grp2* in the *data* disk group.

Example 10-82 Using the chqg command

```
ASMCMD [+] > chqg -G data quota_grp1 quota 200G
```

```
ASMCMD [+] > chqg -G data quota_grp2 quota UNLIMITED
```

10.7.3 lsfg

Purpose

List file groups in a disk group.

Syntax and Description

```
lsfg [-G disk_group] [--filegroup file_group] [--suppressheader]
```

The following table lists the options for the lsfg command.

Table 10-78 Options for the lsfg command

Option	Description
<i>-G disk_group</i>	Specifies the disk group.

Table 10-78 (Cont.) Options for the *lsfg* command

Option	Description
<code>-filegroup file_group</code>	Specifies the file group.
<code>-suppressheader</code>	Suppresses the headers in the output of the command.

If the disk group option is specified, then `lsfg` returns only information about file groups in the specified disk group. If the file group option is specified, then `lsfg` returns only information about that file group.

The information about file groups is retrieved from `V$ASM_FILEGROUP` and `V$ASM_FILEGROUP_PROPERTY` views.

The following example shows the use of the `lsfg` command. The first command in the example displays the properties of the file group `FG1` in the `DATA` disk group. The second command lists all the file groups and their properties in the `DATA` disk group.

Example 10-83 Using the *lsfg* command

```

ASMCMD [+] > lsfg -G DATA --filegroup FG1
Group Number  FileGroup Number  Property  Value  File Type
1              1                 REDUNDANCY  MIRROR
1              1                 PRIORITY   HIGH
1              1                 REDUNDANCY  HIGH    CONTROLFILE

ASMCMD [+] > lsfg -G DATA
Group Number  FileGroup Number  Incarnation  Name                Client Type  Client
Name  QuotaGroup Number  QuotaGroup Incarn
1      0                    1            DEFAULT_FILEGROUP
0      0                    0            0
1      1                    1            FG1                  1
CLIENT1  0                    0
    
```

10.7.4 *lsqg*

Purpose

Lists the quota groups in a disk group.

Syntax and Description

```
lsqg [-G disk_group] [--quotagroup quota_group] [--suppressheader]
```

The following table lists the options for the `lsqg` command.

Table 10-79 Options for the *lsqg* command

Option	Description
<code>-G disk_group</code>	Specifies the disk group.
<code>-quotagroup quota_group</code>	Specifies the quota group.
<code>-suppressheader</code>	Suppresses the headers in the output of the command.

If the disk group option is specified, then `lsqg` returns only information about quota groups in the specified disk group. If the quota group option is specified, then `lsqg` returns only information about that quota group.

The information about file groups is retrieved from `V$ASM_QUOTAGROUP` view.

The following example shows the use of the `lsqg` command. The first command in the example displays the properties of the file group `QG1` in the `DATA` disk group. The second command lists all the quota groups and their properties in the `DATA` disk group. The third command lists all the quota groups with name `QG1` in all the disk groups mounted by the current Oracle ASM instance. The fourth command lists all the quota groups in all the disk groups mounted by the current Oracle ASM instance.

Example 10-84 Using the `lsqg` command

```

ASMCMD [+] > lsqg -G DATA --quotagroup QG1
Used_Quota_MB  Quota_Limit_MB
100             200

ASMCMD [+] > lsqg -G DATA
Quotagroup_Num  Quotagroup_Name  Used_Quota_MB  Quota_Limit_MB
1                GENERIC          12357          262143
2                QG1              100            200
3                QG2              400            420

ASMCMD [+] > lsqg --quotagroup QG1
Group_Num  Quotagroup_Num  Used_Quota_MB  Quota_Limit_MB
1          2          100            200
3          5          821            1024

ASMCMD [+] > lsqg
Group_Num  Quotagroup_Num  Quotagroup_Name  Used_Quota_MB  Quota_Limit_MB
1          1          GENERIC          12357          262143
1          2          QG1              100            200
1          3          QG2              400            420
2          1          GENERIC          54000          262143
3          1          GENERIC          11             262143
3          5          QG1              821            1024
    
```

10.7.5 mkfg

Purpose

Creates a file group in a disk group based on an XML configuration file or script.

Syntax and Description

```
mkfg { config_file.xml | 'contents_of_xml_script' }
```

The following table lists the options for the `mkfg` command.

Table 10-80 Options for the `mkfg` command

Option	Description
<i>config_file.xml</i>	Name of the XML file that contains the configuration information for the file group. <code>mkfg</code> searches for the XML file in the current working directory of ASMCMD unless a path is specified

Table 10-80 (Cont.) Options for the mkfg command

Option	Description
<code>contents_of_xml_script.xml</code>	The contents of an XML script enclosed in single quotes.

The root element of the XML configuration file or script is the `<file_group>` tag and has these attributes:

- `name`: file group name
- `dg`: disk group name to which the file group belongs
- `client_type`: database, cluster, volume
- `client_name`: database, CDB, PDB, cluster or volume name

You can set properties using these tag:

- `name`: property name
- `value`: property value
- `file_type`: property file type (optional)
- `template`: Oracle ASM template name (optional)

The following is an example of an XML configuration file for `mkfg`. The configuration file creates a file group named `filegroup1` on disk group `dg_data1` for a database named `sample`. The file group is associated with quota group `quotagroup1`. The files in this file group have the redundancy set to `mirror`, except for the control file, which has a high redundancy. All other attributes are set to the default value.

```
<filegroup name="filegroup1" dg="dg_data1" client_type="database"
client_name="sample">
  <p name="redundancy" value="mirror"/>
  <p name="redundancy" value="high" file_type="controlfile"/>
  <p name="quota_group" value="quotagroup1"/>
</filegroup>
```

For information about Oracle ASM File Groups, refer to "[Managing Oracle ASM Flex Disk Groups](#) (page 4-72)".

Example

The following examples show the use of the `mkfg` command using an XML configuration file and an XML script.

Example 10-85 Using the mkfg command

```
ASMCMD [+] > mkfg config_file.xml
```

```
ASMCMD [+] > mkfg '<filegroup name="filegroup1" dg="dg_data1" client_type="database"
client_name="sample">
  </filegroup>'
```

10.7.6 mkqg

Purpose

Adds a quota group to the disk group.

Syntax and Description

```
mkqg -G disk_group quota_group [property] [value]
```

The following table lists the options for the `mkqg` command.

Table 10-81 Options for the `mkqg` command

Option	Description
<code>-G <i>disk_group</i></code>	Specifies the disk group.
<code><i>quota_group</i></code>	Specifies the quota group.
<code><i>property</i></code>	Specifies the name of the property.
<code><i>value</i></code>	Specifies the value of the property.

The `mkqg` command adds the specified quota group with an optional the specified property to the specified disk group. For information about Oracle ASM file and quota groups, refer to "[Managing Oracle ASM Flex Disk Groups](#) (page 4-72)".

The following are examples of the `mkqg` command. The first example adds the quota group `quota_grp1` to the data disk group without specifying a property. The second example adds the quota group `quota_grp2` to the data disk group while specifying the quota property.

Example 10-86 Using the `mkqg` command

```
ASMCMD [+] > mkqg -G data quota_grp1
```

```
ASMCMD [+] > mkqg -G data quota_grp2 quota 100G
```

10.7.7 mvfg

Purpose

Moves a file group in a disk group to the specified quota group.

Syntax and Description

```
mvfg -G disk_group --filegroup file_group quota_group
```

The following table lists the options for the `mvfg` command.

Table 10-82 Options for the `mvfg` command

Option	Description
<code>-G <i>disk_group</i></code>	Specifies the disk group.

Table 10-82 (Cont.) Options for the mvfg command

Option	Description
<code>-filegroup file_group</code>	Specifies the file group.
<code>quota_group</code>	Specifies the quota group.

The following is an example of the mvfg command. In the example, the file group FG1 in the DATA disk group is moved to the quota group QG1

Example 10-87 Using the mvfg command

```
ASMCMD [+] > mvfg -G DATA --filegroup FG1 QG1
```

10.7.8 rmfg

Purpose

Removes an existing file group from a disk group.

Syntax and Description

```
rmfg [-r] disk_group file_group
```

The following table lists the options for the rmfg command.

Table 10-83 Options for the rmfg command

Option	Description
<code>-r</code>	Specifies a recursive operation.
<code>disk_group</code>	Specifies the disk group.
<code>file_group</code>	Specifies the file group.

The `-r` option specifies a recursive removal, including contents. All files in the disk group that the file group contains are removed at the same time that the file group is removed. You must specify this option to remove a file group that contains any files. If you omit this option and the file group is not empty, then ASMCMD returns an error and does not remove the file group.

The following example shows the use of the rmfg command. The first command in the example removes file group FG1 from the DATA1 disk group. The second command removes file group FG2, and its contents, from the DATA2 disk group.

Example 10-88 Using the rmfg command

```
ASMCMD [+] > rmfg DATA1 FG1
ASMCMD [+] > rmfg -r DATA2 FG2
```

10.7.9 rmqg

Purpose

Removes a quota group from a disk group.

Syntax and Description

```
rmqg -G disk_group quota_group
```

The following table lists the options for the `rmqg` command.

Table 10-84 Options for the `rmqg` command

Option	Description
<code>-G <i>disk_group</i></code>	Specifies the disk group.
<code><i>quota_group</i></code>	Specifies the quota group.

The following example shows the use of the `rmqg` command. The command in the example removes quota group QG1 from the DATA disk group.

Example 10-89 Using the `rmqg` command

```
ASMCMD [+] > rmqg -G DATA QG1
```

10.8 ASMCMD Template Management Commands

This topic provides a summary of the ASMCMD template management commands.

[Table 10-85](#) (page 10-87) lists the ASMCMD template management commands with brief descriptions.

Table 10-85 Summary of ASMCMD template management commands

Command	Description
chtmpl (page 10-87)	Changes the attributes of a template.
lstmpl (page 10-88)	Lists disk group templates.
mktmpl (page 10-89)	Adds a template to a disk group.
rmtmpl (page 10-90)	Removes a template from a disk group.

10.8.1 chtmpl

Purpose

Changes the attributes of a template.

Syntax and Description

```
chtmpl -G diskgroup { [--striping {coarse|fine}]  
  [--redundancy {high|mirror|unprotected} ] [--primary {hot|cold}]  
  [--secondary {hot|cold}] } template
```

[Table 10-86](#) (page 10-87) lists the syntax options for the `chtmpl` command.

Table 10-86 Options for the `chtmpl` command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the disk group.

Table 10-86 (Cont.) Options for the `chtml` command

Option	Description
<code>template</code>	Name of the template to change.
<code>--striping {coarse fine}</code>	Striping specification, either <code>coarse</code> or <code>fine</code> .
<code>--redundancy {high mirror unprotected}</code>	Redundancy specification, either <code>high</code> , <code>mirror</code> , or <code>unprotected</code> .
<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.

At least one of the following options is required: `--striping`, `--redundancy`, `--primary`, or `--secondary`.

For information about updating templates, see "[Modifying a Disk Group Template](#) (page 5-23)".

Example

The following is an example of the `chtml` 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 10-90 Using the ASMCMD `chtml` command

```
ASMCMD [+] > chhtml -G data --redundancy high --striping fine mytemplate
```

10.8.2 `lstmpl`

Purpose

Lists all templates for all disk groups or the templates for a specified disk group.

Syntax and Description

```
lstmpl [--suppressheader] [-l] [-G diskgroup] [pattern]
```

[Table 10-87](#) (page 10-88) lists the syntax options for the `lstmpl` command.

Table 10-87 Options for the `lstmpl` command

Option	Description
<code>-G <i>diskgroup</i></code>	Specifies disk group name.
<code>--suppressheader</code>	Suppresses column headings.
<code>-l</code>	Displays all details.
<code><i>pattern</i></code>	Displays the templates that match <code>pattern</code> expression.

For information about template attributes, see "[Template Attributes](#) (page 5-20)".

Example

The following is an example of the `lstmpl` command. The example lists all details of the templates in the data disk group.

Example 10-91 Using the ASMCMD `lstmpl` command

```
ASMCMD [+] > lstmpl -l -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	AUDIT_SPILLFILES	COARSE	Y	MIRROR	COLD	COLD
DATA	1	AUTOBACKUP	COARSE	Y	MIRROR	COLD	COLD
DATA	1	AUTOLOGIN_KEY_STORE	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	FLASHFILE	COARSE	Y	MIRROR	COLD	COLD
DATA	1	INCR XTRANSPORT BACKUPSET	COARSE	Y	MIRROR	COLD	COLD
DATA	1	KEY_STORE	COARSE	Y	MIRROR	COLD	COLD
DATA	1	OCRFIL	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	VOTINGFILE	COARSE	Y	MIRROR	COLD	COLD
DATA	1	XTRANSPORT	COARSE	Y	MIRROR	COLD	COLD
DATA	1	XTRANSPORT BACKUPSET	COARSE	Y	MIRROR	COLD	COLD

10.8.3 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 10-88 (page 10-89) lists the syntax options for the `mktmpl` command.

Table 10-88 Options for the `mktmpl` command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the disk group.
<code><i>template</i></code>	Name of the template to create.
<code>--striping {coarse fine}</code>	Striping specification, either <code>coarse</code> or <code>fine</code> .
<code>--redundancy { high mirror unprotected}</code>	Redundancy specification, either <code>high</code> , <code>mirror</code> , or <code>unprotected</code> .

Table 10-88 (Cont.) Options for the *mktmpl* command

Option	Description
<code>--primary { hot cold }</code>	Intelligent Data Placement specification for primary extents, either hot or cold region.
<code>--secondary { hot cold }</code>	Intelligent Data Placement specification for secondary extents, either hot or cold region.

For information about adding templates, see "[Adding Templates to a Disk Group](#) (page 5-22)".

Example

The following is an example of the *mktmpl* command that adds the *mytemplate* template to the *data* disk group. The new template has the redundancy set to mirror and the striping set to coarse.

Example 10-92 Using the ASMCMD *mktmpl* command

```
ASMCMD [+] > mktmpl -G data --redundancy mirror --striping coarse mytemplate
```

10.8.4 *rmtmpl*

Purpose

Removes a template from a disk group.

Syntax and Description

```
rmtmpl -G diskgroup template
```

[Table 10-89](#) (page 10-90) lists the syntax options for the *rmtmpl* command.

Table 10-89 Options for the *rmtmpl* command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the disk group.
<code><i>template</i></code>	Name of the template to delete.

For information about template attributes, see "[Template Attributes](#) (page 5-20)".

Example

The following is an example of the *rmtmpl* command that removes the *mytemplate* template from the *data* disk group.

Example 10-93 Using the ASMCMD *rmtmpl* command

```
ASMCMD [+] > rmtmpl -G data mytemplate
```

10.9 ASMCMD File Access Control Commands

This topic provides a summary of 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](#) (page 4-61).

When administering users with ASMCMD, the users must be existing operating system users and their user names must have corresponding operating system user IDs. However, only users in the same cluster as the Oracle ASM instance can be validated.

Note:

When modifying ownership of files, the previous owner of a file cannot be dropped while the file is still open even after the ownership of the file has changed.

[Table 10-90](#) (page 10-91) lists the ASMCMD Oracle ASM File Access Control commands with brief descriptions.

Table 10-90 Summary of ASMCMD file access control commands

Command	Description
chgrp (page 10-91)	Changes the user group of a file or list of files.
chmod (page 10-92)	Changes permissions of a file or list of files.
chown (page 10-93)	Changes the owner of a file or list of files.
groups (page 10-94)	Lists the user groups to which a user belongs.
grpmod (page 10-95)	Adds or removes users from an existing user group.
lsgrp (page 10-96)	Lists user groups.
lsusr (page 10-96)	Lists users in a disk group.
mkgrp (page 10-97)	Creates a new user group.
mkusr (page 10-98)	Adds a user to disk group.
passwd (page 10-98)	Changes the password of a user.
rmgrp (page 10-99)	Deletes a user group.
rmusr (page 10-99)	Deletes a user from a disk group.
rpusr (page 10-100)	Replaces one a user with another in a disk group.

10.9.1 chgrp

Purpose

Changes the user group of a file or list of files.

Syntax and Description

```
chgrp usergroup file [file ...]
```

[Table 10-91](#) (page 10-92) lists the syntax options for the `chgrp` command.

Table 10-91 Options for the *chgrp* command

Option	Description
<i>usergroup</i>	Name of the user group.
<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 the user must also be either the owner or a member of the group for this command to succeed.

This command accepts a file name or multiple file names separated by spaces.

If you change the user group settings of an open file, then the operation currently running on the file completes using the old user group settings. The new settings take effect when re-authentication is required.

Examples

The following are examples of the *chgrp* command that change the Oracle ASM user group of the specified files.

Example 10-94 Using the ASMCMD *chgrp* command

```
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
```

10.9.2 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.

[Table 10-92](#) (page 10-92) lists the syntax options for the *chmod* command.

Table 10-92 Options for the *chmod* command

Option	Description
6	Read write permissions

Table 10-92 (Cont.) Options for the chmod command

Option	Description
4	Read only permissions
0	No permissions
u	Owner permissions, used with <code>r</code> or <code>w</code>
g	Group permissions, used with <code>r</code> or <code>w</code>
o	Other user permissions, used with <code>r</code> or <code>w</code>
a	All user permissions, used with <code>r</code> or <code>w</code>
+	Add a permission, used with <code>r</code> or <code>w</code>
-	Removes a permission, used with <code>r</code> or <code>w</code>
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.

You can only set file permissions to read-write, read-only, and no permissions. You cannot set file permissions to write-only.

If you change the permission settings of an open file, then the operation currently running on the file completes using the old permission settings. The new settings take effect when re-authentication is required.

To view the permissions on a file, use the ASMCMD `ls` command with the `--permission` option. See "[ls](#) (page 10-47)".

Examples

The following are examples of the `chmod` command that change the permissions of the specified files.

Example 10-95 Using the ASMCMD chmod command

```
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
```

10.9.3 chown

Purpose

Changes the owner of a file or list of files.

Syntax and Description

```
chown user[:usergroup ] file [file ...]
```

[Table 10-93](#) (page 10-94) lists the syntax options for the `chown` command.

Table 10-93 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.

If you change the ownership settings of an open file, then the following applies:

- The operation currently running on the file completes using the old ownership and user group settings. The new settings take effect when re-authentication is required.
- The new owner of the file cannot be dropped from the disk group until all instances in a cluster are restarted in a rolling manner.

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 10-96 Using the ASMCMD `chown` command

```
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
```

10.9.4 groups

Purpose

Lists all the user groups to which the specified user belongs.

Syntax and Description

```
groups diskgroup user
```

[Table 10-94](#) (page 10-94) lists the syntax options for the `groups` command.

Table 10-94 Options for the `groups` command

Option	Description
<i>diskgroup</i>	Name of the disk group to which the user belongs.

Table 10-94 (Cont.) Options for the groups command

Option	Description
<i>user</i>	Name of the user.

Examples

The following is an example of the `groups` command that displays the user groups of the `oracle1` user to which the `oracle1` user belongs.

Example 10-97 Using the ASMCMD groups command

```
ASMCMD [+] > groups data oracle1
asm_data
```

10.9.5 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 10-95](#) (page 10-95) lists the syntax options for the `grpmod` command.

Table 10-95 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 10-98 Using the ASMCMD grpmod command

```
ASMCMD [+] > grpmod --add fra asm_fra oracle1 oracle2

ASMCMD [+] > grpmod --delete data asm_data oracle2
```

10.9.6 lsgrp

Purpose

Lists all Oracle ASM user groups or only groups that match a specified pattern.

Syntax and Description

```
lsgrp [--suppressheader][-a] [ -G diskgroup ] [ pattern ]
```

[Table 10-96](#) (page 10-96) lists the syntax options for the `lsgrp` command.

Table 10-96 Options for the `lsgrp` command

Option	Description
<code>--suppressheader</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 10-99 Using the ASMCMD lsgrp command

```
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
```

10.9.7 lsusr

Purpose

Lists Oracle ASM users in a disk group.

Syntax and Description

```
lsusr [--suppressheader][-a] [-G diskgroup] [ pattern ]
```

[Table 10-97](#) (page 10-97) lists the syntax options for the `lsusr` command.

Table 10-97 Options for the *lsusr* command

Option	Description
-a	List all users and the disk groups to which the users belong.
--suppressheader	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 Id assigned to the user.

Example 10-100 Using the ASMCMD *lsusr* command

```
ASMCMD [+] > lsusr -G data
User_Num OS_ID OS_Name
3        1001  grid
1        1021  oracle1
2        1022  oracle2
```

10.9.8 mkgrp

Purpose

Creates a new Oracle ASM user group.

Syntax and Description

```
mkgrp diskgroup usergroup [user] [user...]
```

[Table 10-98](#) (page 10-97) lists the syntax options for the *mkgrp* command.

Table 10-98 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 10-101 Using the ASMCMD mkgrp command

```
ASMCMD [+] > mkgrp data asm_data oracle1 oracle2
```

10.9.9 mkusr

Purpose

Adds an operating system (OS) user to a disk group.

Syntax and Description

```
mkusr diskgroup user
```

[Table 10-99](#) (page 10-98) lists the syntax options for the `mkusr` command.

Table 10-99 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 to add.

The user to be added must be a valid operating system user. Only a user authenticated as `SYSASM` can run 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 10-102 Using the ASMCMD mkusr command

```
ASMCMD [+] > mkusr data oracle1
```

```
ASMCMD [+] > mkusr fra oracle2
```

10.9.10 passwd

Purpose

Changes the password of a user.

Syntax and Description

```
passwd user
```

[Table 10-100](#) (page 10-98) lists the syntax options for the `passwd` command.

Table 10-100 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 10-103 Using the ASMCMD passwd command

```
ASMCMD [+] > passwd oracle2
Enter old password (optional):
Enter new password: *****
```

10.9.11 rmgrp

Purpose

Removes a user group from a disk group.

Syntax and Description

```
rmgrp diskgroup usergroup
```

[Table 10-101](#) (page 10-99) lists the syntax options for the `rmgrp` command.

Table 10-101 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 user group. To ensure that those files have a valid group, explicitly update those files to a valid user group. See "[chgrp](#) (page 10-91)".

The command must be run by the owner of the user 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 10-104 Using the ASMCMD rmgrp command

```
ASMCMD [+] > rmgrp data asm_data
```

10.9.12 rmusr

Purpose

Deletes an operating system (OS) user from a disk group.

Syntax and Description

```
rmusr [-r] diskgroup user
```

[Table 10-102](#) (page 10-100) lists the syntax options for the `rmusr` command.

Table 10-102 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.
<code>diskgroup</code>	Specifies the name of the disk group from which the user is to be deleted.
<code>user</code>	Name of the user to delete.

The command requires the SYSASM privilege to run.

Examples

The following is an example of the `rmusr` command that removes the `oracle2` user from the `data` disk group.

Example 10-105 Using the ASMCMD `rmusr` command

```
ASMCMD [+] > rmusr data oracle2
```

10.9.13 `rpusr`

Purpose

Replaces one operating system (OS) user with another in a disk group.

Syntax and Description

```
rpusr diskgroup user1 user2
```

[Table 10-103](#) (page 10-100) lists the syntax options for the `rpusr` command.

Table 10-103 Options for the `rpusr` command

Option	Description
<code>diskgroup</code>	Specifies the name of the disk group.
<code>user1</code>	Name of the user to be replaced.
<code>user2</code>	Name of user that replaces the existing user in the disk group.

The `rpusr` command replaces one operating system user with another in an Oracle ASM disk group. The replacement user must not be a user currently in the disk group user list.

If the command succeeds, all files that were previously owned by current user are now owned by the replacement user. The current user is automatically removed from the user list of the disk group. Adding the current user back to the disk group is possible but this user will not own any files.

When changing the ownership of an open file, the new owner cannot be dropped with the owner's files until all instances within a cluster are restarted.

The command requires the SYSASM privilege to run.

Examples

The following is an example of the `rpusr` command that replaces the `oracle1` user with the `oracle2` user in the data disk group.

Example 10-106 Using the ASMCMD rpusr command

```
ASMCMD [+] > rpusr data oracle1 oracle2
```


Part III

Oracle Automatic Storage Management Cluster File System

Part III discusses Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM), in the following chapters:

- [Introducing Oracle ACFS and Oracle ADVM](#) (page 11-1)
- [Using Views to Display Oracle ACFS Information](#) (page 12-1)
- [Administering Oracle ACFS with Oracle Enterprise Manager](#) (page 13-1)
- [Managing Oracle ACFS and Oracle ADVM With ASMCA](#) (page 14-1)
- [Managing Oracle ADVM with ASMCMD](#) (page 15-1)
- [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1)

Introducing Oracle ACFS and Oracle ADVM

Oracle ASM Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) provide key components of storage management.

This chapter describes the components of Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM).

This chapter provides concepts and an overview of Oracle ACFS and Oracle ADVM features with the following topics:

- [Overview of Oracle ACFS](#) (page 11-1)
- [Understanding Oracle ACFS Concepts](#) (page 11-4)
- [Understanding Oracle ACFS Administration](#) (page 11-15)
- [Overview of Oracle ASM Dynamic Volume Manager](#) (page 11-49)

See Also:

- ["Managing Oracle ADVM Volumes in a Disk Group](#) (page 4-21)" for information about using SQL*Plus to administer volumes
 - ["Features Enabled By Disk Group Compatibility Attribute Settings](#) (page 4-57)" for information about Oracle ACFS features and disk group attribute settings
 - [Managing Oracle ASM With ASMCA](#) (page 9-1) for information about using Oracle ASM Configuration Assistant to administer Oracle ACFS
 - [Using Views to Display Oracle ACFS Information](#) (page 12-1) for information about using views to obtain Oracle ACFS information
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about volume management commands
 - [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1) for information about Oracle ACFS operating system utilities
-

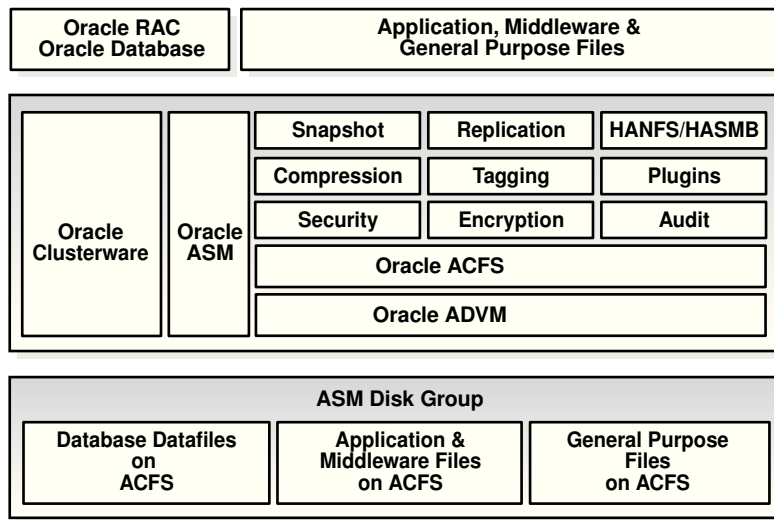
11.1 Overview of Oracle ACFS

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 all customer files.

Oracle ACFS supports Oracle Database files and application files, including executables, database data files, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and all other general-purpose application file data. Oracle ACFS conforms to POSIX standards for Linux and UNIX, and to Windows standards for Windows.

An Oracle ACFS file system communicates with Oracle ASM and is configured with Oracle ASM storage, as shown in the following figure.

Figure 11-1 Oracle ACFS Storage Layers



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

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 (standalone) and cluster configurations.

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 also be managed with Oracle ASM Configuration Assistant. 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). However, CIFS clients on Windows cannot use ACLs when interfacing with Oracle ACFS Linux, Solaris, or AIX servers, but can use ACLs with Oracle ACFS on Windows.

In addition to sharing file data, Oracle ACFS provides additional storage management services including support for the Oracle Grid Infrastructure clusterwide mount registry, dynamic online file system resizing, and multiple space efficient snapshots for each file system.

Oracle ACFS contributes to the overall Oracle storage management by providing:

- A general-purpose standalone server and cluster file system solution that is integrated with Oracle ASM and Oracle Clusterware 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
- Standalone and clusterwide shared Oracle Database homes, all Oracle Database files, and application data
- Uniform, coherent shared file access and clusterwide naming of all customer application files
- Integration with Oracle Clusterware High Availability Resources

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 millions of files.

Oracle ACFS provides support for sparse files. Oracle ACFS sparse files greatly benefit NFS client write operations which are commonly received out of order by the NFS server and the associated Oracle ACFS file system. Usually when an application writes beyond the end of file, storage is allocated and zeroes inserted beyond the old end of file and the beginning of the new. With this feature, a hole remains in the file instead of the inserted zeroes. Oracle ACFS then fills these holes with zeroes in memory when the holes are read. The sparse files feature benefits NFS performance and also the performance and disk utilization of other applications that intentionally perform this type of writing. In addition, there are also storage savings benefits for files that are inherently sparse, meaning they have a lot of unused space, such as some image files for virtual machines. For sparse file support, the `COMPATIBLE .ADVM` disk group attribute must be set to 12.2 or greater.

Oracle ACFS file systems are generally mounted on all Oracle Cluster Synchronization Services (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.

The following list provides important information about Oracle ACFS:

- For all applications, Oracle ACFS performance is best with larger `write()` sizes, such as 8 K or larger.
- For best performance, use the Deadline I/O Scheduler for the disks in the disk group on a Linux system.

- When creating Oracle ACFS file systems on Windows, log on as a Windows domain user. Also, when creating files in an Oracle ACFS file system on Windows, you should be logged in as a Windows domain user to ensure that the files are accessible by all nodes.
- When using a file system across cluster nodes on Windows platforms, the best practice is to mount the file system using a domain user, to ensure that the security identifier is the same across cluster nodes. Windows security identifiers, which are used in defining access rights to files and directories, use information which identifies the user. Local users are only known in the context of the local node. Oracle ACFS uses this information during the first file system mount to set the default access rights to the file system.

Oracle ACFS does not support any files associated with the management of Oracle ASM, such as files in the Oracle Grid Infrastructure home and in the Oracle ASM diagnostic directory.

- Oracle ACFS does not support Oracle Cluster Registry (OCR) and voting files.
- Oracle ACFS functionality requires that the disk group compatibility attributes for ASM and ADVM be set to 11.2 or higher.
- To use an Oracle ACFS file system for an Oracle Database home, the release level must be Oracle 11g Release 2 (11.2) or later.

See Also:

- [Overview of Oracle ASM Dynamic Volume Manager](#) (page 11-49) for an overview of Oracle ADVM
- [Oracle Clusterware Resources and Oracle ACFS Administration](#) (page 11-18) for information Oracle Clusterware High Availability Resources
- [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
- [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1) and [Using Views to Display Oracle ACFS Information](#) (page 12-1) for information about managing and monitoring Oracle ACFS
- *Oracle Database Installation Guide* for the Windows platform for information about Oracle Base permissions when a file system is mounted under Oracle Base.

11.2 Understanding Oracle ACFS Concepts

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

- [About Oracle ACFS](#) (page 11-5)
- [About the Oracle ACFS Mount Model and Namespace](#) (page 11-5)
- [About Oracle ACFS and Database Data Files](#) (page 11-6)
- [About Oracle ACFS and Oracle Database Homes](#) (page 11-8)

- [About Oracle ASM Dynamic Volume Manager](#) (page 11-10)
- [About the Oracle ACFS Driver Model](#) (page 11-10)
- [About the Oracle ACFS Mount Registry](#) (page 11-11)
- [About Oracle ACFS Snapshots](#) (page 11-11)
- [About Oracle ACFS and Backup and Restore](#) (page 11-14)
- [About Oracle ACFS Integration with Oracle ASM](#) (page 11-15)
- [About Oracle ACFS and External Tables on Windows](#) (page 11-15)

11.2.1 About Oracle ACFS

Oracle ACFS is designed as a general-purpose, standalone server and clusterwide file system that delivers support for all customer files. 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 Oracle ASM Configuration Assistant (ASMCA).

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 and 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 with 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 file 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 standalone server applications.

Oracle ACFS presents single system file access semantics across cluster configurations. Applications and users on all cluster members are always presented with the same view of shared Oracle ACFS file data, supported by the Oracle ACFS clusterwide user and metadata cache coherency mechanism.

11.2.2 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 leaves. The namespace design is a single-file system naming model for both standalone server and cluster configurations. This design enables each cluster member to present shared files to cluster applications using the same path names, simplifying multi-node application and user access, and 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.

It is best practice for Oracle ACFS file systems to be Oracle Clusterware managed with Oracle Clusterware resources to ensure they are properly handled during Oracle Grid Infrastructure startup and shutdown.

You can explicitly use the `mount` command. However, if the resource has been created, then the file system may already be mounted.

11.2.3 About Oracle ACFS and Database Data Files

Note:

The best practice for data files in Oracle ACFS is to use a clusterwide resource for the Oracle ACFS File system that stores the data files.

If the data files are added after the database installation is complete, you must modify the database dependencies to list the new Oracle ACFS file system. Failure to do so results in errors with application reliability. To specify the file systems in the dependency list, use the `SRVCTL` database object command to modify the Oracle ACFS paths used by the resource.

Oracle ACFS in Oracle Grid 12c supports all database files starting with Oracle Database 11g Release 2 (11.2.0.4), except for data files and redo logs in an Oracle Restart (standalone server) configuration. Oracle ACFS can be configured for use with the database particularly to leverage Oracle ACFS snapshots for database testing and development. To support database files, the `COMPATIBLE .ADVM` attribute must be set to `12.1` or higher for the disk group that contains the Oracle ACFS file system.

Support for database data files on Windows begins with Oracle Grid 12c Release 1 (12.1.0.2). For support of database files on Windows, the `COMPATIBLE .ADVM` attribute must be set to `12.1.0.2` or higher.

Support for database data files on Oracle Exadata (Linux) begins with Oracle Grid 12c Release 1 (12.1.0.2). However, Oracle ACFS does not currently have the ability to push database operations directly into storage.

Oracle ACFS additionally supports all database files for Oracle Database 10g Release 2 (10.2.0.4 and 10.2.0.5) on Oracle Exadata (Linux) storage only. For database file support with Oracle Database 10g Release 2 (10.2.0.4 and 10.2.0.5) on Oracle Exadata storage, the following conditions must be met:

- When creating an Oracle Database with DBCA, you must set the `REMOTE_LISTENER` initialization parameter to `your_scan_vip:1521` otherwise DBCA fails during the create process.
- You must modify all the start and stop dependencies of the database instance resources to ensure that the resources start when starting Oracle Clusterware.

The following list provides important information about using Oracle ACFS with database files:

- Oracle ACFS support includes all file types supported by Oracle ASM.
- Oracle ACFS does not support data files or redo logs in an Oracle Restart configuration.
- When storing database data files on Oracle ACFS, you must set the `FILESYSTEMIO_OPTIONS` initialization parameter to `setall`; other settings are not supported. To achieve optimal performance with database data files, set `ASM` and `ADVM` compatibility attributes to `12.1` or higher for the disk group that contains the Oracle `ADVM` volume intended to hold the data files. For volumes created before 12.1.0.2, set the stripe columns to `1`, or set the stripe columns to `8`.

and the stripe width to 1 MB. Volumes created while running 12.1.0.2 or higher already default to the high performance configuration (stripe columns = 8 and stripe width = 1 MB).

- To obtain optimal database performance with snapshots, the snapshots must be created after the `ADVM` compatibility attribute is set to 12.1 or higher.
- Use a 4 K or larger database block size and tablespace block size with Oracle ACFS for best performance.
- For best performance when using Oracle Exadata, use a 4 K metadata file system.
- Oracle ACFS does not support databases with 2 K block sizes.
- Do not export Oracle ACFS file systems containing database files through NFS to access them from a DNFS client. This configuration is not supported.
- If a data file is configured to automatically extend, then the size of the increments should be large enough to ensure that the extend operation occurs infrequently. Frequent automatic extends have a negative performance impact.
- Running a workload in a snapshot reduces resources for the primary workload running on the base files because the storage is shared between the base file system and the snapshots. To run test scenarios in Oracle ACFS snapshots without impacting the primary workload, copy the file system and then run test workloads on snapshots created in the copied file system.
- Using Oracle ACFS replication or encryption with database files on Oracle ACFS is not supported. Instead, consider Oracle Data Guard or Oracle GoldenGate. Oracle GoldenGate is an Oracle product sold independently of the Oracle Database. To encrypt database data files on Oracle ACFS, Oracle recommends Oracle Advanced Security. Oracle Advanced Security provides Transparent Data Encryption (TDE) to encrypt data files for entire tablespaces.
- Databases that share storage with snapshots or with the base of the file system can become fragmented under active online transaction processing (OLTP) workloads. This fragmentation can cause the location of the data in the volume to be discontinuous for sequential scans. Oracle ACFS automatically defragments these files in the background. Fragmentation is reported through the `acfsutil defrag dir` and `file` commands, and it can also be viewed with the `acfsutil info file` command. The `acfsutil defrag dir` and `file` commands also enable on-demand defragmentation of a file in the event the automatic defragmentation does not occur quickly enough. Defragmentation may also be needed when database data files reside in file systems enabled with Oracle ACFS compression.

See Also:

- [Creating Oracle ACFS Resources](#) (page 11-19) for information about Oracle ACFS resources
 - [What Types of Files Does Oracle ASM Support?](#) (page 5-1) for a list of file types supported by Oracle ASM
 - [volcreate](#) (page 15-2) for information about creating a volume
 - [mkfs](#) (page 16-28) for information about setting the metadata block size with `mkfs -i`
 - [Oracle ACFS Command-Line Utilities](#) (page 16-143) for information about `acfsutil defrag` and `acfsutil info` commands
 - *Oracle Database Reference* for information about the `REMOTE_LISTENER` initialization parameter
 - *Oracle Clusterware Administration and Deployment Guide* for information about resource dependencies
 - *Oracle Data Guard Concepts and Administration* for information about other replication options for database files on Oracle ACFS
 - *Oracle Database Advanced Security Guide* for information about Transparent Data Encryption (TDE)
-

11.2.4 About Oracle ACFS and Oracle Database Homes

An Oracle ACFS file system can be configured as an Oracle Database home.

When installing a database in a cluster, it is required to use a shared Oracle ACFS file system for the database home. You can use an Oracle ACFS file system for an Oracle Database home with Oracle 11g Release 2 (11.2) or later.

When installing Oracle Software, there must be a separate Oracle base (`ORACLE_BASE`) associated with each operating system user. For example, there should be a separate Oracle base for a grid user and a database user.

You can locate the Oracle Database base (`ORACLE_BASE` for database) directory and home (`ORACLE_HOME` for database) directory on an Oracle ACFS file system. 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).

The Oracle Grid Infrastructure base (`ORACLE_BASE` for grid) directory and home (`ORACLE_HOME` for grid) directory cannot be located on the Oracle ACFS file system because the Oracle ACFS file system cannot be created until Oracle Grid Infrastructure is installed.

One or more Oracle Database homes on Oracle ACFS can be created under the same mount point with each home using a separate Oracle ACFS file system.

After the installation of Oracle 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 also use the Oracle ASM Configuration Assistant (ASMCA) or Oracle ACFS commands to create the file system.

Note:

When an Oracle ACFS file system contains an Oracle Database home or Oracle Database uses the file system for any file storage, the file system must have an Oracle ACFS file system resource. If you have not used Oracle ASM Configuration Assistant to setup the mount point, then you must use Server Control Utility (SRVCTL) commands to set up Oracle Database dependencies.

In an Oracle Grid Infrastructure clusterware configuration, a clusterwide Oracle ACFS resource is required when using Oracle ACFS for the database home. To enable the database owner to act on the resource, the owner must be specified as a permitted user when creating the resource. You can specify the database owner with the `-u` option of the `srvctl add filesystem` or `acfsutil registry` command. Root privilege is required when adding the resource in Linux or Unix environments.

After the Oracle ACFS file system and resources are 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 Oracle Universal Installer (OUI) Database Software installation.

You can use the `srvctl start filesystem` command to manually mount the Oracle ACFS file system.

Note:

When additional Oracle ACFS file systems are added to an Oracle Database after creation, they must be specified in the Oracle Database Resource Dependency list. Failure to do so results in errors with application reliability. To specify the file systems in the dependency list, use the SRVCTL database object command to modify the Oracle ACFS paths used by the resource.

Oracle ACFS file systems can be also configured for use as a home for applications. 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.

To reduce contention on an Oracle ACFS file system in an Oracle RAC environment where the Oracle Database home is shared on Oracle ACFS, Oracle Database auditing operating system files should be configured as node specific. For a node-specific setup, you must ensure that the `AUDIT_FILE_DEST` initialization parameter in the configuration file of each database instance points to a unique location rather than one location for all the database instances.

For example, if you have a database with the Oracle name set to `TEST` and you want to ensure that the location of `AUDIT_FILE_DEST` initialization parameter for each database instance, such as `TEST1` or `TEST2`, points to a node specific location for that instance, you can run the following SQL statement:

```
SQL> ALTER SYSTEM SET AUDIT_FILE_DEST='$ORACLE_BASE/admin/adump/TEST/@'
SCOPE=SPFILE SID='*';
```

In the previous example, @ expands to the ORACLE_SID of each instance. If ORACLE_BASE has been set to /acfsmounts in this example, then that value could have been used in place of the ORACLE_BASE variable.

See Also:

- [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1) for information about using Oracle ACFS commands to create a file system
 - [Creating an Oracle ACFS File System for Database Use](#) (page 14-5) for information about using ASMCA
 - [Creating Oracle ACFS Resources](#) (page 11-19) for additional information Oracle ACFS resources
 - *Oracle Real Application Clusters Administration and Deployment Guide* for information about SRVCTL commands
 - *Oracle Database Installation Guide* for your environment for information about the installation of Oracle Database software and storage options and about Optimal Flexible Architecture (OFA) recommendations for Oracle base and home directories
 - *Oracle Grid Infrastructure Installation Guide* for your environment for information about the installation of Oracle Grid Infrastructure software and storage options
 - *Oracle Database Administrator's Guide* for information about using Oracle Managed files
 - *Oracle Database Reference* for information about the AUDIT_FILE_DEST initialization parameter
-

11.2.5 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 send 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, refer to "[Overview of Oracle ASM Dynamic Volume Manager](#) (page 11-49)".

11.2.6 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. Refer to "[Understanding Oracle ACFS I/O Failure Console Messages](#) (page 17-29)".

11.2.7 About the Oracle ACFS Mount Registry

The Oracle ACFS mount registry supports Oracle Grid Infrastructure cluster configurations, but does not support Oracle Restart configurations.

File systems that are to be mounted persistently (across restarts) 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, in Oracle Restart configurations the automatic mounting of registered Oracle ACFS file systems is not supported.

By default, an Oracle ACFS file system that is inserted into the Oracle ACFS mount registry is automatically mounted on all cluster members, including cluster members that are added after the registry addition. However, the Oracle ACFS mount registry also accommodates standalone 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 mount registry is implemented using Oracle Clusterware resources, specifically the Oracle ACFS resource. The Oracle ACFS resource actions are designed to automatically mount a file system only one time for each Oracle Grid Infrastructure initialization to avoid potential conflicts with administrative actions to dismount a given file system.

See Also:

- [Oracle ACFS and Oracle Restart](#) (page 17-6) for more information about Oracle ACFS and Oracle Restart
 - [acfsutil registry](#) (page 16-160) for information about registering an Oracle ACFS file system using the `acfsutil` command
 - [Oracle Clusterware Resources and Oracle ACFS Administration](#) (page 11-18) for more information about the implementation of the mount registry and Oracle ACFS resources
 - *Oracle Real Application Clusters Administration and Deployment Guide* for information about the `SRVCTL srvctl add filesystem` command and registering an Oracle ACFS file system
 - *Oracle Clusterware Administration and Deployment Guide* for information about Oracle Clusterware resources
-
-

11.2.8 About Oracle ACFS Snapshots

An Oracle ACFS snapshot is an online, read-only or read-write, 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. The snapshots are created in the `.ACFS/snaps/` directory of the file system. They are always online while the file system is mounted. Consequently, an Oracle ACFS snapshot can support the online recovery of files inadvertently modified or deleted

from a file system. 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.

Up to 1023 read-only, read-write, or combination of read-only and read-write snapshot views are supported for each file system, enabling flexible online file recovery solutions spanning multiple views to be employed. On 64-bit systems, Oracle ACFS supports 1023 snapshots. The total number of snapshots can be any combination of read-only and read-write snapshots. To support the creation of more than 63 snapshots, the disk group compatibility attribute for ADVM must be set to 12.1.0.2 or greater. Also, the following conditions must be met to create more than 63 snapshots.

- All snapshots of the file system that were created with `COMPATIBLE.ADVM` set to less than 12.1.0.2 must be deleted.
- All snapshots of the file system that were created after setting `COMPATIBLE.ADVM` to 12.1.0.2 or greater while older snapshots existed must be deleted.

Oracle ACFS read-write snapshots enable fast creation of an snapshot image that can be both read and written without impacting the state of the Oracle ACFS file system hosting the snapshot images. You can use read-write snapshots for:

- Testing of new versions of application software on production file data reflected in the read-write snapshot image without modifying the original production file system
- Running test scenarios on a real data set without modifying the original production file system

To use Oracle ACFS read-write snapshots, the disk group compatibility attribute for ADVM must be set to 11.2.0.3.0 or higher. If you create a read-write snapshot on an existing Oracle ACFS file system from a version earlier than 11.2.0.3.0, then the file system is updated to the 11.2.0.3.0 or higher format. After a file system has been updated to a higher version, an Oracle ACFS file system cannot be reverted to an earlier version, and accordingly cannot be mounted on an earlier Oracle Grid Infrastructure version.

You can create a snapshot from an existing snapshot in the same Oracle ACFS file system. In addition, you can convert a snapshot between read-only and read-write formats. To create from an existing snapshot or convert a snapshot, the disk group compatibility attribute for ADVM must be set to 12.1 or higher. In addition, creation from an existing snapshot is not permitted if there are:

- Any snapshots present in the file system that were created with the ADVM compatibility set to less than 12.1
- Any snapshots of the file system that were created after ADVM compatibility was set to 12.1 but while 11.2 snapshots existed

Oracle ACFS snapshot storage is maintained within the file system, eliminating the management of separate storage pools for file systems and snapshots. Oracle ACFS file systems can be dynamically resized to accommodate additional file and snapshot storage requirements.

You cannot modify security or encryption metadata in read-write snapshots except for enabling or disabling security or encryption. No other alteration is permitted on Oracle ACFS security or encryption metadata in a snapshot. If a file was not secured by a security realm in the snapshot, it cannot be realm secured by adding the corresponding file in the active file system to a security realm. If a file was not

encrypted in the snapshot, that file cannot be encrypted by encrypting the corresponding file in the active file system.

A new file created in a realm-secured directory in a read-write snapshot inherits the realm security attributes of the parent directory. If the realm protecting the new file has encryption turned on, the file is encrypted with the encryption parameters set in the realm. If the realm protecting the new file has encryption turned off, the file is decrypted. Files and directories in a read-write snapshot cannot be added to or removed from any security realm.

Files in a read-write snapshot can be encrypted, decrypted, or rekeyed if the operation target is a path specified for a file or directory of the read-write snapshot. However, if an encryption, decryption, or rekey operation is specified at the file system level, then the operation does not process files and directories of snapshots in the `.ACFS/snaps/` directory.

All Oracle ACFS snapshot operations are serialized clusterwide in the kernel. For example, if a snapshot create operation is initiated at the same time as a snapshot delete operation, then both operations would complete, but they would not run in parallel inside of the kernel. One operation would complete before the other was started.

The `acfsutil snap duplicate` commands can be invoked manually to enable a target file system or writable snapshot to track the evolution of a source file system or snapshot with minimal overhead. Oracle ACFS updates the target after each `apply` operation to indicate its current contents, and checks before starting the next `apply` operation to ensure that the starting point of the incoming snapshot duplication stream matches the endpoint of the last-applied stream. The target remains writable by any and all applications even when it is in use by `acfsutil snap duplicate apply`. Oracle ACFS determines any updates to the target by a process other than the `apply` process, and if such an update occurs, invalidates the internal metadata that track the target contents. This prevents further `apply` operations to the target, and avoids the possibility of a future `apply` operation yielding unexpected results.

The use of the `acfsutil snap duplicate` commands interacts with Oracle ACFS replication operations. A target file system in use by `acfsutil snap duplicate apply` cannot be specified for use as a file system for replication. And a file system in use by replication cannot be specified as the target for `acfsutil snap duplicate apply`.

To run the `acfsutil snap duplicate` commands, the disk group compatibility attribute for ADVM must be set to at least 12.2.

Additional `acfsutil snap` commands are available to manage duplicate snapshots and remaster an Oracle ADVM volume with an existing snapshot.

Oracle ACFS snapshots are administered with the `acfsutil snap` commands.

Note:

The `link()` and `rename()` system calls fail if an attempt is made to link or rename a file in the Oracle ACFS file system and a file in any associated read-write snapshot, or vice versa. Any tools which use the `link()` and `rename()` system calls, such as `ln` and `mv`, also fail in the same scenario.

See Also:

- [acfsutil snap create](#) (page 16-127) for additional information on the target used with the `acfsutil snap duplicate` command
 - [Oracle ACFS Replication](#) (page 11-36) for information about Oracle ACFS replication
 - [Oracle ACFS Security](#) (page 11-22) for information about Oracle ACFS security
 - [Oracle ACFS Encryption](#) (page 11-25) for information about Oracle ACFS encryption
 - [Oracle ACFS Command-Line Tools for Snapshots](#) (page 16-125) for information about the `acfsutil snap` commands
-
-

11.2.9 About Oracle ACFS and Backup and Restore

Oracle ACFS runs on operating system platforms as a native file system technology supporting native operating system file system application programming interfaces (APIs). Consequently, backup applications that access files using the native operating system file system interfaces are able to access and backup Oracle ACFS file systems and 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.

Note:

When Oracle ACFS is used with Oracle Recovery Manager (RMAN) backups or restore, you should tune RMAN to improve performance because the default values may not be ideal for Oracle ACFS file systems. For example, you can adjust RMAN parameters to match the buffer size with the volume stripe width or Oracle ASM AU size. Also, you can adjust the number of buffers to match the number of devices in the Oracle ASM disk group.

See Also:

- [Oracle ACFS Tagging](#) (page 11-39) for information about using common operating system utilities to preserve Extend Attributes for tagging definitions
 - [md_backup](#) (page 10-66) and [md_restore](#) (page 10-67) for information about backing up and restoring Oracle ACFS metadata
 - [My Oracle Support](#) (<https://support.oracle.com>) for articles about tuning RMAN tuning, such as *RMAN Performance Tuning Using Buffer Memory Parameters* (Doc ID 1072545.1) and *RMAN Performance Troubleshooting* (Doc ID 1326686.1)
-
-

11.2.10 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 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. However, I/O does not go through Oracle ASM nor through the Oracle ASM proxy, but rather goes directly to the underlying Oracle ASM storage.

For information about Oracle ACFS and Oracle ASM operations, refer to "[Oracle ACFS and Dismount or Shutdown Operations](#) (page 11-22)".

11.2.11 About Oracle ACFS and External Tables on Windows

To access an external table stored on an Oracle ACFS file system on Windows, the external table must be created with the `DISABLE_DIRECTORY_LINK_CHECK` access parameter.

See Also:

- *Oracle Database Utilities* for information about the `DISABLE_DIRECTORY_LINK_CHECK` access parameter
 - *Oracle Database SQL Language Reference* for information about creating external tables
-
-

11.3 Understanding Oracle ACFS Administration

This section describes Oracle ACFS administration and contains the following topics:

- [Oracle ACFS and File Access and Administration Security](#) (page 11-16)
- [Oracle ACFS and Grid Infrastructure Installation](#) (page 11-17)
- [Oracle ACFS Configuration](#) (page 11-17)
- [Oracle Clusterware Resources and Oracle ACFS Administration](#) (page 11-18)
- [Oracle ACFS and Dismount or Shutdown Operations](#) (page 11-22)
- [Oracle ACFS Security](#) (page 11-22)
- [Oracle ACFS Encryption](#) (page 11-25)
- [Oracle ACFS Compression](#) (page 11-28)
- [Oracle ACFS Auditing](#) (page 11-29)
- [Oracle ACFS Replication](#) (page 11-36)
- [Oracle ACFS Tagging](#) (page 11-39)
- [Using Replication with Auditing, Encryption, and Security](#) (page 11-41)
- [Oracle ACFS Plugins](#) (page 11-43)
- [Oracle ACFS Accelerator Volume](#) (page 11-43)
- [Oracle ACFS NAS Maximum Availability eXtensions](#) (page 11-44)

11.3.1 Oracle ACFS and File Access and Administration Security

Oracle ACFS supports both traditional Unix-style file access control classes (user, group, other) for Linux environments 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 environments 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.

Additional fine grain access control is provided for Oracle ACFS file systems with the security infrastructure feature.

See Also:

- [Oracle ACFS Security](#) (page 11-22) for information about Oracle ACFS security infrastructure
 - [Oracle ACFS Encryption](#) (page 11-25) for information about Oracle ACFS encryption
 - [About Privileges for Oracle ASM](#) (page 3-28) for information about Oracle ASM privileges
 - [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1) for information about administering Oracle ACFS
-
-

11.3.2 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.

11.3.3 Oracle ACFS Configuration

After a Grid Infrastructure installation and with an operational Oracle Clusterware, you can use Oracle ASM Configuration Assistant (ASMCA) 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 and ASMCMD command line tools. File systems can be created using operating system command-line tools.

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.

See Also:

- [Basic Steps to Manage Oracle ACFS Systems](#) (page 16-2) for an example of the specific actions required to create a file system
 - [ASMCA GUI Tool for Managing Oracle ACFS and Oracle ADVM](#) (page 14-1) for information about managing Oracle ACFS file systems with ASMCA
 - [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1) for information about using Oracle ACFS commands to create a file system
 - [Oracle Clusterware Resources and Oracle ACFS Administration](#) (page 11-18) for information about configuring Oracle Clusterware resources for use with the database
-
-

11.3.4 Oracle Clusterware Resources and Oracle ACFS Administration

Oracle Clusterware resources support all aspects of Oracle ACFS. The resources are responsible for enabling and disabling volumes, loading drivers and mounting and unmounting file systems.

This section discusses the following topics:

- [Summary of Oracle ACFS Resource-based Management](#) (page 11-18)
- [High Availability Actions](#) (page 11-19)
- [Creating Oracle ACFS Resources](#) (page 11-19)
- [Node-Local or Clusterwide File Systems](#) (page 11-20)
- [Monitoring Oracle ACFS resources](#) (page 11-21)
- [Stopping Oracle ACFS resources](#) (page 11-21)
- [Oracle ACFS resource Limitations](#) (page 11-21)

11.3.4.1 Summary of Oracle ACFS Resource-based Management

The following list provides a summary of 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 ADVN 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 additional information, see "[Oracle ACFS Drivers Resource Management](#) (page 17-4)" and "[Oracle ACFS Driver Commands](#) (page 17-6)".

- When a volume is created, Oracle ADVN creates a resource with the name of `ora.DISKGROUP.VOLUME.advn`. This resource is usually managed through transparent high availability calls from Oracle ASM and requires no user interaction. However, the user may choose to use the SRVCTL command interface to start and stop volumes as well as control the default state of the volume after an Oracle ASM restart. This is especially beneficial in a large cluster or an Oracle Flex ASM cluster, as volumes on other nodes may be operated upon.

In addition, these Oracle ADVN resources can be used by other resources in the Oracle Clusterware stack to maintain dependency chains. Dependency chains ensure that the resources a program requires to run are available. For instance, if a resource was monitoring a backup application that was backing up to Oracle

ADVM volume, the backup application would want to ensure that it specified the Oracle ADVM volume resource in its `START` and `STOP` dependency list. Because the Oracle ADVM volume resource will enable the volume, this ensures that the volume is available before the backup begins.

- Oracle ACFS file systems are either manually mounted or dismounted using an Oracle ACFS or Oracle Clusterware command-line tool, or automatically mounted or dismounted based on an Oracle Clusterware resource action.

For example, a file system hosting an Oracle Database home is 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.

Oracle ACFS file system resources provide the following actions:

- `MOUNT`

During the `START` operation the resource mounts the file system on the path configured in the resource. The Oracle ACFS file system resource requires all components of the Oracle ASM stack to be active (volume device, ASM) and ensures that they are active before attempting the mount.

- `UNMOUNT`

During the `STOP` operation, the resource attempts to unmount a file system.

- Oracle provides two resource types for Oracle Highly Available NFS. For more information, refer to "[High Availability Network File Storage for Oracle Grid Infrastructure](#) (page 11-44)".

As with all Oracle Clusterware resources, these resources provide for high availability by monitoring the underlying device, file system, or driver to ensure that the object remains available. In the event that the underlying object becomes unavailable, each resource attempts to make the underlying object available again.

11.3.4.2 High Availability Actions

The following are the actions of the High Availability resources:

- Oracle ACFS resource

This resource attempts to unmount the file system. After the unmount has succeeded, the resource remounts the file system, making the file system available again. If processes are active on the file system during unmount, the resource identifies and terminates those processes.

- Oracle ADVM resource

This resource attempts to disable any volume device, and then reenables the volume device. At that point, any configured Oracle ACFS resource can remount the file system. If processes are active on the volume during this period, the resource identifies and terminates the processes.

11.3.4.3 Creating Oracle ACFS Resources

Oracle ACFS resources can be created with the following methods:

- Oracle ASM Configuration Assistant (ASMCA) provides a GUI that exposes the most common functionality. In all cases, creating a file system resource does not

format the underlying file system. Attempts to start the resource require the user to format the file system either manually or with ASMCA.

- SRVCTL provides a highly flexible command line utility for creating Oracle ACFS file system resources through the *filesystem* object. Oracle ACFS resources created through this mechanism have access to the full feature set, including server pools.
- `acfsutil` commands provide an alternative method to create Oracle ACFS file system resources using the *registry* object. Oracle ACFS resources created through this methodology have access to a limited set of options.

The differences between SRVCTL and `acfsutil` commands are:

- Oracle ACFS resources created through SRVCTL and specifying a server pool or list of nodes are only mounted on one of those nodes. (node-local)
- Oracle ACFS resources created through SRVCTL can take advantage of Oracle Server Pools.
- Oracle ACFS resources created through `acfsutil` commands and specifying a list of nodes are mounted on all listed nodes. (node-local)
- Oracle ACFS resources created through `acfsutil` commands are created with `AUTOSTART` set to `ALWAYS`.
- Oracle ACFS resources created through SRVCTL allow for advanced Application ID functionality. Using this functionality enables the resource type to be set by the administrator. After the type is set, other resources can depend on this type, allowing different node-local file systems to be used to fulfill dependencies on each node. In a simplified example, this would allow the administrator to have a different device mounted on the `/log` directory on each node of the cluster, and be able to run an Apache resource. The Apache resource would specify the new type in its resource dependency structure, rather than specifying an individual resource.
- Oracle ACFS resources created through SRVCTL can specify additional `AUTOSTART` parameters. These parameters can be used to prevent the resource from starting on stack startup, to always force the resource to start, or to only start the resource if it was previously running.
- Oracle ACFS resources created through SRVCTL have access to functionality such as accelerator volumes.

The common elements of both SRVCTL and `acfsutil` commands are:

- User
This is an additional user that can act upon the resource. By default, you must be the `root` user to start and stop an Oracle ACFS resource.
- Options
These are mount options that should be used to mount the file system when the resource is starting.

11.3.4.4 Node-Local or Clusterwide File Systems

When creating Oracle ACFS file system resources, you can create a node-local file system or to create a clusterwide file system.

- **Node-local**
This file system type is limited to the number of nodes it can mount on. Depending on if it is created with `SRVCTL` or `acfsutil` commands, it may only mount on one node, a subset of nodes, or all the configured nodes. In some cases, this could look the same as a full cluster configuration, but if new nodes are added to the cluster, the file system is not automatically mounted on them without modifying the list of allowable nodes.
- **Clusterwide**
This type of file system mounts on all nodes of the cluster, with no exceptions. When new members are added to the cluster, the file system is automatically available on them. This type of resource is required for certain configurations, such as Oracle Database or Oracle HANFS.

11.3.4.5 Monitoring Oracle ACFS resources

Similar to all Oracle Clusterware resources, Oracle ACFS resources enables you to monitor the state of the system. You can do this monitoring with the following commands:

- **Using SRVCTL commands**
When the command `srvctl status filesystem` or `srvctl status volume` is run, the output of the command reports if the file system is mounted or the volume is enabled, and which nodes this is true on.
- **Using CRSCCTL commands**
When the `crsctl status resource` command is run, a state of `ONLINE` is reported for each resource that is available, whether through a mounted file system or an enabled volume. A state of `OFFLINE` is reported for each resource that is not available, whether through an unmounted file system or a disabled volume. Additional status may be presented in the `STATUS` field of this output.

11.3.4.6 Stopping Oracle ACFS resources

You can stop Oracle ACFS file system resources with the following methods:

- You can stop the entire Oracle Clusterware stack. When the Oracle Clusterware stack is stopped, all Oracle ACFS resources are automatically stopped.
- To stop individual resources, you can use `SRVCTL` management commands with the `filesystem` or `volume` object. The command may require the `-force` option if there are other resources that are depending on the resource that you are attempting to stop.
- You may engage a manual action, such as running `umount` on a file system or by manually stopping a volume using `ASMCMD` or `SQL*Plus` commands. In this case, the Oracle ACFS resource transitions to the `OFFLINE` state automatically.

11.3.4.7 Oracle ACFS resource Limitations

Oracle ACFS has the following resource limitations:

- All Oracle ACFS resources require `root` privileges to create.
- All Oracle ACFS resources require `root` privileges to remove.

- All Oracle ACFS file system resources require `root` privileges to act upon, such as starting and stopping the resources, but can be configured to allow another user, such as a database user, to do so. In this case, the `root` user must be used to configure the resource.
- All Oracle ADVM volume resources allow the `ASMADMIN` user to act upon them.
- All Oracle ACFS resources are only available in Oracle RAC mode. Oracle ACFS resources are not supported in Oracle Restart configurations. For more information about Oracle ACFS and Oracle Restart, refer to "[Oracle ACFS and Oracle Restart](#) (page 17-6)".

11.3.5 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 run 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 "[About Shutting Down an Oracle ASM Instance](#) (page 3-22)". For information about dismounting a disk group, see "[Mounting and Dismounting Disk Groups](#) (page 4-67)".

11.3.6 Oracle ACFS Security

Oracle ACFS security provides realm-based security for Oracle ACFS file systems, enabling you to create realms to specify security policies for users and groups to determine access on file system objects.

This security feature provides a finer-grained access control on top of the access control provided by the operating system. Oracle ACFS security can use the encryption feature to protect the contents of realm-secured files stored in Oracle ACFS file systems.

Oracle ACFS security uses realms, rules, rule sets, and command rules to enforce security policies.

- An Oracle ACFS security realm is a group of files or directories that are secured for access by a user or a group of users. Realms are defined with rule sets which contain groups of rules that apply fine grain access control. Oracle ACFS security realms can also be used as containers to enable encryption.
- Oracle ACFS security rules are Boolean expressions that evaluate to true or false based on a system parameter on which the rule is based.
- Oracle ACFS rule sets are collection of rules. Rule sets evaluate to TRUE or FALSE based on the evaluation of the rules a rule set contains.
- Oracle ACFS command rules are associations of the file system operation to a rule set. For example, the association of a file system create, delete, or rename operation to a rule set. Command rules are associated with an Oracle ACFS realm.

An existing operating system user must be designated as the first Oracle ACFS security administrator and an existing operating system group must be designated as the security administrator *admin* group. Security administrators must be members of the designated security group. Additional users can be designated as security administrators. An Oracle ACFS security administrator can manage encryption for an Oracle ACFS file system on a per-realm basis. An Oracle ACFS security administrator is authenticated for security operations with a security realm password, not the operating system password of the user.

The first security administrator is created during the initialization of Oracle ACFS security with the `acfsutil sec init` command which is run by the root user. When the first security administrator is created, the administrator is assigned a password that can be changed by the administrator. Each time a security administrator runs an `acfsutil sec` command, the administrator is prompted for the security password. The security realm passwords for administrators are stored in a wallet created during the security initialization process. This wallet is located in the Oracle Cluster Registry (OCR).

Auditing and diagnostic data are logged for Oracle ACFS security. The log files include information such as `acfsutil` commands that have been run, the use of security or system administrator privileges, and run-time failures such as realm check authorization failures.

Auditing events, such as realm creation or encryption enabled, are written to these log files only if auditing is not enabled for on the file system. If auditing is enabled, these events are written into the audit trail. Diagnostic messages related to security and encryption are always written to the `sec-hostname_fsid.log` file regardless of whether auditing is enabled or not.

Logs are written to the following files:

- `mount_point/.Security/realm/logs/sec-hostname_fsid.log`
The directory is created with `acfsutil sec prepare` command and protected by Oracle ACFS security.
- `GRID_HOME/log/hostname/acfs/security/acfssec.log`
The messages that are logged to this file are for commands that are not associated with a specific file system, such as `acfsutil sec init`. The directory is created during installation and is owned by the root user.

When an active log file grows to a pre-defined maximum size (10 MB), the file is automatically moved to `log_file_name.bak`, the administrator is notified, and logging continues to the regular log file name. When the administrator is notified, the

administrator must archive and remove the `log_file_name.bak` file. If an active log file grows to the maximum size and the `log_file_name.bak` file exists, logging stops until the backup file is removed. After the backup log file is removed, logging restarts automatically.

Oracle ACFS security protects the following objects from unauthorized accesses:

- Realm-secured directories and user files
The directories and files reside on a file system secured by Oracle ACFS security.
- The Oracle ACFS security directory (`mount_point/.Security`) and its contents
The security directory contains the log files in plain-text format and a security metadata backup file in XML format. The log files generated by Oracle ACFS security can only be accessed by valid Oracle ACFS security administrators.
- Oracle ACFS security objects
These objects are the security realms, rules, and rule sets used to manage Oracle ACFS security.

Access to files in a security realm of an Oracle ACFS file system must be authorized by both the security realm and the underlying operating system permissions, such as (owner, group, other) permissions on Linux and Access Control Lists (ACLs) on Windows. Each access to a realm-secured file is first checked for security realm authorization. If the access is authorized by the security realm, then access to the files is checked by the underlying operating system access control checks. If both checks pass, access is allowed to the realm-secured file.

Note the following when working with Oracle ACFS security:

- Oracle ACFS security does not provide any protection for data sent on the network.
- A copy of a realm-protected file is not realm-protected unless the copy is made in a security realm-protected directory.

Some applications, such as the `vi` editor, re-create a file when the file is modified. The modified file is saved as a temporary file, the original file is removed, and temporary file is copied with the original file name as the destination name. This process creates a new file. If the new file is created in a realm-protected directory, the security policies of the realm also apply to the new file. If the new file is not created in a realm-protected directory, then the new file is not realm-protected. If you are planning to copy a realm-protected file, you should ensure that the parent directory is also security realm protected.

Security policies also apply to any temporary files created in a realm-protected directory.

To use Oracle ACFS security functionality on Linux, the disk group compatibility attributes for ASM and ADVM must be set to 11.2.0.2 or higher. To use Oracle ACFS security functionality on Windows, the disk group compatibility attributes for ASM and ADVM must be set to 11.2.0.3 or higher. .

Security information for Oracle ACFS file systems is displayed in the `V $ASM_ACFS_SECURITY_INFO` view.

See Also:

- [Oracle ACFS Auditing](#) (page 11-29) for information about Oracle ACFS auditing
 - [acfsutil sec prepare](#) (page 16-96) for information about preparing Oracle ACFS file systems
 - [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - [About Oracle ACFS Snapshots](#) (page 11-11) for information about Oracle ACFS security and snapshots
 - [Using Views to Display Oracle ACFS Information](#) (page 12-1) for information about V\$ASM_ACFS views
 - [Securing Oracle ACFS File Systems](#) (page 16-9) and [Oracle ACFS Command-Line Tools for Security](#) (page 16-86) for information about the `acfsutil sec` command-line functions to configure security for Oracle ACFS file systems
 - [Managing Security and Encryption for Oracle ACFS with ASMCA](#) (page 14-4) for information about using ASMCA to configure Oracle ACFS security
 - Your operating system-specific (OS) documentation for information about setting up OS users and OS groups
-

11.3.7 Oracle ACFS Encryption

Oracle ACFS encryption enables you to encrypt data stored on disk (data-at-rest).

The encryption feature protects data in an Oracle ACFS file system in encrypted format to prevent unauthorized use of data in the case of data loss or theft. Both encrypted and non-encrypted files can exist in the same Oracle ACFS file system.

Some encryption functionality requires system administrator privileges. This functionality includes the commands for initiating, setting, and reconfiguring encryption.

System administrators and Oracle ACFS security administrators can initiate encryption operations. Also, unprivileged users can initiate encryption for files they own.

Oracle ACFS encryption provides two type of encryption keys:

- **File Encryption Key**
This is a key for a file and is used to encrypt the data in the file.
- **Volume Encryption Key**
This is a key for a file system and is used to encrypt the file encryption keys.

You must first create the encryption key store, then specify file system-level encryption parameters and identify the directories. No extra steps are required for a user to read encrypted files if the user has the appropriate privileges for accessing the file data.

Oracle ACFS encryption supports both Oracle Cluster Registry (OCR) and Oracle Key Vault as a key store. Both OCR and Oracle Key Vault can be used in the same cluster. However, a single file system uses either OCR or Oracle Key Vault as a key store, but not both. Oracle Key Vault is currently only available with file systems on Linux.

If you are using OCR as a key store, you should back up the OCR after creating or updating an encryption key to ensure there is an OCR backup that contains all of the volume encryption keys (VEKs) for the file system.

Oracle ACFS encryption protects data stored on secondary storage against the threat of theft or direct access to the storage medium. Data is never written to secondary storage in plaintext. Even if physical storage is stolen, the data stored cannot be accessed without the encryption keys. The encryption keys are never stored in plaintext. The keys are either obfuscated, or encrypted using a user-supplied password.

An Oracle ACFS security administrator can manage encryption parameters on a per-realm basis. After a file is placed under realm security, file-level encryption operations are not allowed on that file. Even if the realm security allows the file owner or the root user to open the file, file-level encryption operations are blocked. Encryption of realm-protected files is managed entirely by the Oracle ACFS security administrator, who can enable and disable encryption for files at a security realm level.

After a directory has been added to a security realm, all files created in the directory inherit the realm-level encryption parameters, not the directory or file system-level parameters. When a file is removed from its last security realm, the file is encrypted or decrypted to match the file system-level encryption status. The file is not re-encrypted to match file system-level parameters if it has been encrypted with security realm parameters.

A system administrator cannot rekey realm-secured files at the file system or file level. To ensure all realm-secured files are encrypted with the most recent volume encryption key (VEK), you must first remove encryption from all realms, and then re-enable encryption. This action re-encrypts all files with the most recent VEK.

Auditing and diagnostic data are logged for Oracle ACFS encryption. The log files include information such as `acfsutil` commands that have been run, the use of security or system administrator privileges, and run-time failures. Logs are written to the following files:

- `mount_point/.Security/encryption/logs/encr-hostname_fsid.log`
The directory is created with `acfsutil encr set` command and protected by Oracle ACFS security if security is enabled.
- `GRID_HOME/log/hostname/acfs/security/acfssec.log`
The messages that are logged to this file are for commands that are not associated with a specific file system, such as `acfsutil encr init`. The directory is created during installation and is owned by the root user.

When an active log file grows to a pre-defined maximum size (10 MB), the file is automatically moved to `log_file_name.bak`, the administrator is notified, and logging continues to the regular log file name. When the administrator is notified, the administrator must archive and remove the `log_file_name.bak` file. If an active log file grows to the maximum size and the `log_file_name.bak` file exists, logging stops until the backup file is removed. After the backup log file is removed, logging restarts automatically.

Note the following when working with Oracle ACFS encryption:

- A copy of an encrypted file is not encrypted unless the copy of the file is made in an encrypted directory.

Some applications, such as the `vi` editor, re-create a file when the file is modified. The modified file is saved as a temporary file, the original file is removed, and temporary file is copied with the original file name as the destination name. This process creates a new file. The new file is not encrypted unless it is created in an encrypted directory. If you are planning to copy an encrypted file, you should ensure that the parent directory is also encrypted.

- Using encryption with database files on Oracle ACFS is not supported.
- Oracle ACFS encryption cannot be used with password-protected (PKCS) wallets if any of the file systems using encryption are configured to be mounted with the Oracle ACFS mount registry.
- The `acfsutil encr on`, `acfsutil encr off`, and `acfsutil encr rekey` commands are not supported on files greater than 128 megabytes. The operation of those commands on a file greater than 128 megabytes fails and the resulting error message displays alternative actions.

To use Oracle ACFS encryption functionality on Linux, the disk group compatibility attributes for ASM and ADVM must be set to 11.2.0.2 or higher. The disk group compatibility attributes for ASM and ADVM must be set to 11.2.0.3 or higher on Linux for the following cases:

- If encryption is configured for the first time on Oracle ASM 11g Release 2 (11.2.0.3).
- If encryption parameters must be changed or a new volume encryption key must be created following a software upgrade to Oracle ASM 11g Release 2 (11.2.0.3).

To use Oracle ACFS encryption functionality on Windows, the disk group compatibility attributes for ASM and ADVM must be set to 11.2.0.3 or higher.

Encryption information for Oracle ACFS file systems is displayed in the `V $ASM_ACFS_ENCRYPTION_INFO` view. To configure encryption and manage encrypted Oracle ACFS file systems, you can use the `acfsutil encr` command-line functions and Oracle ASM Configuration Assistant.

See Also:

- [acfsutil encr set](#) (page 16-123) and [acfsutil encr rekey](#) (page 16-122) for information about changing or creating a volume encryption key
 - [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - [About Oracle ACFS Snapshots](#) (page 11-11) for information about Oracle ACFS encryption and snapshots
 - [Using Views to Display Oracle ACFS Information](#) (page 12-1) for information about V\$ASM_ACFS views
 - [Encrypting Oracle ACFS File Systems](#) (page 16-14) and [Oracle ACFS Command-Line Tools for Encryption](#) (page 16-117) for information about managing encryption with Oracle ACFS command-line tools
 - [Managing Security and Encryption for Oracle ACFS with ASMCA](#) (page 14-4) for information about using Oracle ASMCA with encryption features
 - *Oracle Key Vault Administrator's Guide* for information about Oracle Key Vault
-

11.3.8 Oracle ACFS Compression

Oracle ACFS compression is enabled on a specified Oracle ACFS file system for database data files, RMAN files, archive logs, data pump files, and general purpose files.

Redo logs, flashback logs, and control files are not compressed. For database related files, the Oracle Database release must be 11.2.0.4 or greater.

Direct IO as done for database files is done inline and the database file is compressed immediately when written. Cached IO compression is performed asynchronously, after the application has written to the file.

When enabling compression on a file system, existing files are not compressed, only newly-created files are compressed. When compression is disabled, compressed files are not uncompressed.

Compressed files are associated with a compression unit and the compression algorithm operates on this unit. For database files, the size of the unit is made equal to the database block size for maximum efficiency. For non-database files, the unit size is currently 32 kilobytes. When copying data files to an Oracle ACFS file system that has been enabled for compression, use the `acfsutil compress copy` command to ensure that the proper compression unit size is maintained for correct functionality. `lzo` is the default compression algorithm and the only compression algorithm currently supported.

The `acfsutil compress` command sets and resets the compression state of a file system with `acfsutil compress on` and `acfsutil compress off`. To display the compression state and effectiveness of the compression operation, use the `acfsutil compress info file` command. The `acfsutil info fs` and `acfsutil info file` commands have been enhanced to report on Oracle ACFS compression status.

Compressed files consume less disk space than non-compressed files. However, for applications using the file, the size reported is equal to the uncompressed file size, not the smaller compressed size. Some utilities, such as `ls -l`, report the uncompressed size of the file. Utilities such as `du`, `acfsutil compress info`, and `acfsutil info file`, report the actual disk allocation of the compressed file.

Note the following about Oracle ACFS compression.

- Loopback mounts are not supported with compressed files. If a loopback device is associated with a compressed file, read and write operations to the loopback device fail.
- A loopback device can be associated with an uncompressed file on an Oracle ACFS file system that has been enabled for compression.
- Databases with block sizes of 2 kilobytes or 4 kilobytes are not supported on compression enabled file systems.
- Compression should not be used on file systems that primarily have small files. Small files on compressed file systems consume the default compression unit size of 32 kilobytes.
- Database files in read-write (RW) snapshots of compressed file systems are not defragged automatically. If you want to defrag the database files to improve sequential access, then use the `acfsutil defrag` command. However, defragging may result in the consumption of more storage.
- For Oracle Grid Infrastructure 12c release 2 (12.2.0.1), Oracle ACFS compression is supported on Linux and AIX.
- Oracle ACFS compression is only supported with Oracle ACFS snapshot-based replication that is available starting with Oracle Grid Infrastructure 12c release 2 (12.2.0.1).
- ADVM disk group compatibility must be set to 12.2 or higher.

See Also:

- [Oracle ACFS Command-Line Tools for Compression](#) (page 16-140) for information about the Oracle ACFS compression commands
 - [Oracle ACFS Command-Line Utilities](#) (page 16-143) for information about the `acfsutil info` commands
-
-

11.3.9 Oracle ACFS Auditing

Oracle ACFS auditing provides auditing capabilities for Oracle ACFS security and encryption. This auditing framework produces a separate audit trail for each Oracle ACFS file system on each individual node, and enforces separation of duties regarding the management and review of this audit source.

Audit sources are the source of events, such as Oracle ACFS security and Oracle ACFS encryption. Audit trails are the logs where the audit records are written.

This section contains the following topics:

- [About Oracle ACFS Auditing](#) (page 11-30)

- [Audit Trail File](#) (page 11-31)
- [File Access Events](#) (page 11-33)
- [Privilege Use Events](#) (page 11-34)
- [Evaluation Result Events](#) (page 11-35)

11.3.9.1 About Oracle ACFS Auditing

Both Oracle ACFS security and encryption are also audit sources, and these sources can be enabled and disabled by an Oracle ACFS audit manager. These sources generate events as a result of the execution of Oracle ACFS security or encryption commands.

The Oracle ACFS security administrator can enable auditing at the realm level so that security violations and authorizations can also be audited as well as enabling auditing on security to audit all the events executed by a security administrator. An Oracle ACFS security source must be enabled before Oracle ACFS realm security auditing can be used.

Setting the realm auditing policy to audit all authorizations and violations for all command rules can cause the audit trail to quickly increase to its maximum size. Administrators should carefully adjust the auditing level to their requirements and be aware that auditing policies generating more verbose auditing output require additional active monitoring and management, such as archiving and purging, of the audit trail and audit trail backup files.

Along with the generation of a file system audit source, Oracle ACFS auditing allows fine-grained auditing policies to be set separately on each realm basis. The Oracle ACFS auditing capability provides the infrastructure for an audit vault collector to import data into Oracle Audit Vault and Database Firewall. The collector is separate from Oracle ACFS and functions as means for Oracle ACFS auditing data to be imported into Audit Vault Server.

The responsibilities for configuration and management of the audit source are separated into the Oracle ACFS audit manager and Oracle ACFS auditor roles. The system administrator has the authority to add and remove users to and from the Oracle ACFS audit manager and Oracle ACFS auditor operating system (OS) groups.

The Oracle ACFS audit managers have access to the contents of audit sources and can read audit data; however, the audit managers cannot modify the audit sources. The set of Oracle ACFS audit managers is the same across a cluster.

The Oracle ACFS auditors are responsible for viewing and analyzing the contents of the audit source, such as indicating to the Oracle ACFS audit managers which records have been analyzed and archived and are safe to purge. The Oracle ACFS auditors should be the only users on the system with access to the contents of the audit source. The Oracle ACFS auditor do not have the required permissions to remove or purge audit records. The set of Oracle ACFS auditors is the same across a cluster.

The audit archiving process renames audit trail log files (.log) to a audit trail backup file (.log.bak) and generates an XML file, which can be imported by Audit Vault Server. Audit Vault Server has only read access to the audit trail directory and functions as an auditor in this case. After the data from the XML file is imported in the Audit Vault Server, the auditor function marks the audit trail backup file as read, and then audit manager can execute a purge to remove audit trail backup files and XML files.

To configure auditing for an Oracle ACFS file system, run the `acfsutil audit init` command to initialize auditing for Oracle ACFS and then run `acfsutil audit enable` to enable auditing for Oracle ACFS encryption or security on the specified file system.

See Also:

- [Oracle ACFS Command-Line Tools for Auditing](#) (page 16-81) for information about the `acfsutil audit` commands
 - [Oracle ACFS Command-Line Tools for Security](#) (page 16-86) for information about enabling or disabling auditing for specific commands in an Oracle ACFS security realm with the `acfsutil sec realm audit enable` and `acfsutil sec realm audit disable` commands
 - [Views Containing Oracle ACFS Information](#) (page 12-1) for information about views that are relevant to Oracle ACFS auditing
 - *Oracle Audit Vault and Database Firewall Administrator's Guide* for information about the Audit Vault Server
 - Your operating system-specific (OS) documentation for information about setting up OS users and OS groups
-
-

11.3.9.2 Audit Trail File

Audit trail files consist of a set of audit records. Each audit record represents a single event. Audit trail files are located in the `mount_point/.Security/audit` directory.

Audit trail files generated by Oracle ACFS auditing are meant to be available for the following:

- Manual review by an Oracle ACFS auditor using text viewing tools
- Import into Oracle Audit Vault and Database Firewall
- Third party products that can parse and import the audit sources

The audit trail file consists of audit records. There are several different types of audit records, each of which represent a unique type of event and contain different information relevant to diagnosing the event. The types of events are:

- [File Access Events](#) (page 11-33)
- [Privilege Use Events](#) (page 11-34)
- [Evaluation Result Events](#) (page 11-35)

The combination of audit record fields entered in the audit trail file depends on the event type.

Each record is written to the audit trail file as a set of field names and values. Depending on the type of record, the number and type of fields may vary. Fields consist of a name and value pair, in the form `field name:value`, followed by an end of line character.

The audit record fields that can be present in the audit trail file are described in the following list. The string in parenthesis is the field name that appears in the audit trail log file.

- **Timestamp (Timestamp):** The time at which the event occurred, always specified in UTC. The format for the time stamp is: MM/DD/YYYY HH:MM:SS UTC
- **Event Code (Event):** A code identifying the type of event. For the list of evaluation result codes, refer to "[File Access Events](#) (page 11-33)" and "[Privilege Use Events](#) (page 11-34)".
- **Source (Source):** Oracle ACFS
- **User identification (User):** The user who triggered the event. On Linux platforms this is a user ID and on Windows this is the user SID.
- **Group identification (Group):** The primary group of the user who triggered the event. On Linux platforms this is the ID the primary group of the user and on Windows this is the SID of the primary group of the user.
- **Process identification (Process):** The current process ID.
- **Host name (Host):** The host which recorded the event.
- **Application name (Application):** The application name for the current process.
- **Realm name (Realm):** The name of the realm which was violated, or the realm that is authorized and is protecting the file.
- **File name (File):** The file name which the user was accessing.
- **Evaluation Result (Evaluation Result):** This field contains the information about the result of the command executed. For the list of evaluation result codes, refer to "[Evaluation Result Events](#) (page 11-35)".
- **File system Id (FileSystem-ID):**
- **Message (Message):** The message field has the information about the command executed and its result.

[Example 11-1](#) (page 11-32) shows an example of an audit trail file.

Example 11-1 Sample audit trail file

```
Timestamp: 06/08/12 11:00:37:616 UTC
Event: ACFS_AUDIT_READ_OP
Source: Oracle_ACFS
User: 0
Group: 0
Process: 1234
Host: slc01hug
Application: cat
Realm: MedicalDataRealm
File: f2.txt
Evaluation Result: ACFS_AUDIT_REALM_VIOLATION
FileSystem-ID: 1079529531
Message: Realm authorization failed for file ops READ

Timestamp: 06/08/12 11:00:37:616 UTC
Event: ACFS_AUDIT_WRITE_OP
Source: Oracle_ACFS
```

```

User: 102
Group: 102
Process: 4567
Host: slc01hug
Application: vi
Realm: PayrollRealm,SecuredFiles
File: f2.txt
Evaluation Result: ACFS_AUDIT_REALM_AUTH
FileSystem-ID: 1079529531
Message: Realm authorization succeeded for file ops WRITE

```

```

Timestamp: 06/08/12 10:42:20:977 UTC
Event: ACFS_SEC_PREPARE
Source: Oracle_ACFS
User: 507867
Group: 8500
Process: 603
Host: slc01hug
Application: acfsutil.bin
Evaluation Result: ACFS_CMD_SUCCESS
FileSystem-ID: 1079529531
Message: acfsutil sec prepare: ACFS-10627: Mount point '/mnt' is now
prepared for security operations.

```

11.3.9.3 File Access Events

File access events include both realm authorization and violation records. These events share a similar structure with all events, but have a different event code. The Evaluation Result (Evaluation Result) field can contain either ACFS_AUDIT_REALM_VIOLATION or ACFS_AUDIT_REALM_AUTH.

The possible event code (Event) for file access events include the following:

- ACFS_AUDIT_APPENDFILE_OP
- ACFS_AUDIT_CHGRP_OP
- ACFS_AUDIT_CHMOD_OP
- ACFS_AUDIT_CHOWN_OP
- ACFS_AUDIT_CREATEFILE_OP
- ACFS_AUDIT_DELETEFILE_OP
- ACFS_AUDIT_EXTEND_OP
- ACFS_AUDIT_GET_EXTATTR_OP
- ACFS_AUDIT_LINKFILE_OP
- ACFS_AUDIT_MKDIR_OP
- ACFS_AUDIT_MMAPREAD_OP
- ACFS_AUDIT_MMAPWRITE_OP
- ACFS_AUDIT_MUTABLE_OP
- ACFS_AUDIT_OPENFILE_OP
- ACFS_AUDIT_OVERWRITE_OP

- ACFS_AUDIT_READ_OP
- ACFS_AUDIT_READDIR_OP
- ACFS_AUDIT_RENAME_OP
- ACFS_AUDIT_RMDIR_OP
- ACFS_AUDIT_SET_EXTATTR_OP
- ACFS_AUDIT_SYMLINK_OP
- ACFS_AUDIT_TRUNCATE_OP
- ACFS_AUDIT_WRITE_OP

11.3.9.4 Privilege Use Events

Privilege use events include security commands run by the security administrator or system administrator, and encryption commands run by the system administrator or file owners.

The ACFS_AUDIT_INIT, ACFS_SEC_INIT, and ACFS_ENCR_INIT events are written into the global log that is located in Oracle Grid Infrastructure home.

The possible event code (Event) for privilege use events include the following:

- ACFS_AUDIT_ARCHIVE
- ACFS_AUDIT_DISABLE
- ACFS_AUDIT_ENABLE
- ACFS_AUDIT_INIT
- ACFS_AUDIT_PURGE
- ACFS_AUDIT_READ
- ACFS_ENCR_FILE_OFF
- ACFS_ENCR_FILE_ON
- ACFS_ENCR_FILE_REKEY
- ACFS_ENCR_FS_OFF
- ACFS_ENCR_FS_ON
- ACFS_ENCR_INIT
- ACFS_ENCR_SET
- ACFS_ENCR_SET_UNDO
- ACFS_ENCR_VOL_REKEY
- ACFS_ENCR_WALLET_STORE
- ACFS_REALM_AUDIT_DISABLE
- ACFS_REALM_EDIT_ENCR

- ACFS_REALM_AUDIT_ENABLE
- ACFS_SEC_LOAD
- ACFS_SEC_PREPARE
- ACFS_SEC_PREPARE_UNDO
- ACFS_SEC_REALM_ADD
- ACFS_SEC_REALM_CLONE
- ACFS_SEC_REALM_CREATE
- ACFS_SEC_REALM_DELETE
- ACFS_SEC_REALM_DESTROY
- ACFS_SEC_RULE_CREATE
- ACFS_SEC_RULE_DESTROY
- ACFS_SEC_RULE_EDIT
- ACFS_SEC_RULESET_CREATE
- ACFS_SEC_RULESET_DESTROY
- ACFS_SEC_RULESET_EDIT
- ACFS_SEC_SAVE

11.3.9.5 Evaluation Result Events

Evaluation result event codes provide information about the execution status of a command.

The evaluation result event codes can be one of the following:

- ACFS_AUDIT_REALM_VIOLATION – The user executing the command does not have the proper realm access permission to execute the command.
- ACFS_AUDIT_REALM_AUTH - Indicates the result of a realm evaluation.
- ACFS_AUDIT_MGR_PRIV – Audit manager privileges are required, but have not been granted to the user.
- ACFS_AUDITOR_PRIV – Auditor privileges are required, but have not been granted to the user.
- ACFS_CMD_SUCCESS - The command has been successful in performing the task.
- ACFS_CMD_FAILURE - The command has failed in performing the task.
- ACFS_ENCR_WALLET_AUTH_FAIL – A system administrator provides an incorrect password when opening an encryption wallet.
- ACFS_INSUFFICIENT_PRIV – Either file owner or system administrator privileges are required, but have not been granted to the user.
- ACFS_SEC_ADMIN_PRIV – Security administrator privileges are required, but the user is not a security administrator

- `ACFS_SEC_ADMIN_AUTH_FAIL` – A valid security administrator fails to authenticate properly using their Oracle ACFS security administration password
- `ACFS_SYS_ADMIN_PRIV` – System administrator privileges are required, but have not been granted to the user.

11.3.10 Oracle ACFS Replication

Oracle ACFS snapshot-based replication enables replication of Oracle ACFS file systems across a network to a remote site, providing disaster recovery capability for the file system.

The source Oracle ACFS file system of an Oracle ACFS replication relationship is referred to as a primary file system. The target Oracle ACFS file system of an Oracle ACFS replication relationship is referred to as a standby file system.

Note:

- Oracle ACFS replication functionality supports only one standby file system for each primary file system.
- The standby file system is read-only for as long as replication is active on it. Read-write snapshots may be created of the standby if desired.
- A primary site running Linux, Solaris or AIX can replicate to a standby site running any of those operating systems. A primary site running Windows can replicate only to a standby site running Windows.
- The primary and standby sites should be running the same version of the Oracle Grid Infrastructure software. When upgrading the sites, update the standby site first.
- Using replication with database files on Oracle ACFS is not supported.
- Oracle ACFS replication is not supported with Oracle Restart.
- An Oracle Key Vault keystore is not supported on a standby file system with replication.
- Oracle ACFS encryption cannot be undone on a primary file system with replication.

You cannot undo encryption on a file system having active snapshots. If you want to undo encryption on primary file system with active replication, then first terminate replication. After replication has stopped, then undo encryption and start replication again.

A site can host both primary and standby file systems. For example, if there are cluster sites A and B, a primary file system hosted at site A can be replicated to a standby file system at site B. Also, a primary file system hosted at site B can be replicated to a standby file system at site A. However, an Oracle ACFS file system cannot be used simultaneously as a primary and a standby file system.

Oracle ACFS snapshot-based replication operates by recording snapshots of the primary file system. After the initial snapshot is transferred to the standby file system, replication continues by transferring the changes between successive snapshots of the primary to the standby file system. These replication operations can occur either in

constant mode (enabling a new operation to start as soon as the previous one completes), or can be scheduled to occur at fixed intervals. This replication solution is by nature asynchronous.

Oracle ACFS replication uses snapshot functionality on the primary site initially to externalize both the contents of the initial snapshot, and later the differences between two specified snapshots. The result is called a snapshot duplication stream. The replication process then uses snapshot functionality on the standby site to apply this stream to the standby file system, creating a duplicate of the primary file system.

On the primary, because replication works by comparing successive snapshots, it is critical that there be enough disk space available on the site hosting the primary file system to contain the version of the file system recorded in each snapshot, as well as the current file system contents. In addition, it must always be possible to create the snapshots required. Each replication snapshot is deleted when no longer needed.

On the standby, a backup snapshot is created at the end of each replication operation. This snapshot records the latest consistent contents of the standby, and can be used to recover those contents if a permanent outage occurs during the current replication operation. Each backup snapshot is deleted when the following replication operation is complete, so it must always be possible to create a backup snapshot. In addition, enough space must exist for the version of the standby captured in the snapshot and the current file system contents.

You should ensure that the primary and standby file systems do not run out of disk space. If either file system runs out of available storage, you should either expand the file system or, on the primary file system or in read-write snapshots on the standby, remove files from the file system to free up space. You can also configure automatic resize to avoid running out of space.

If the primary file system runs out of space and you decide to free up space by removing files, then you should only remove files that are not being replicated. Replicated files have been stored in a snapshot pending transfer to the standby file system and are not deleted. You can delete any Oracle ACFS snapshots not created by replication.

Oracle ACFS replication uses the `ssh` utility as the transport between the primary and standby clusters. To enable the automated use of `ssh`, replication requires two kinds of *keys* to be configured. These keys must be available on each node where replication is enabled to run.

- On each node in the primary cluster, the system administrator user (the user `root` on non-Windows systems or local `SYSTEM` on Windows) must have a *host key* stored for each node in the standby cluster.
- On each node of the standby cluster, a designated unprivileged user, the *apply user*, must have a public key stored for `root`, or local `SYSTEM`, that is authorized to log in as the *apply user* on that node.

Note that `ssh` is not provided natively on Windows. For more information about the needed keys and how to configure them, and about installing and configuring `ssh` on Windows, refer to [Configuring Oracle ACFS Snapshot-Based Replication](#) (page 18-1).

Before using replication on a file system, ensure that you have checked the following:

- There is sufficient network bandwidth to support replication between the primary and standby file systems.

- The configuration of the sites hosting the primary and standby file systems enable the standby file system to keep up with the rate of change on the primary file system.
- Host keys and user keys for `ssh` have been configured as described previously.

Directories and files in an Oracle ACFS file system can be tagged to select specific objects that you want to replicate in a file system.

Before replicating an Oracle ACFS file system, a replication configuration must be established that identifies information such as the site hosting the primary file system, the site hosting the standby file system, the file system to be replicated, the mount point of the file system, and a list of tags if desired.

The primary and standby sites must share the same user and group configurations, including all `uids` and `gids` in use in the file system. The *apply user* described previously must be configured on each standby node where replication is enabled. This user should be a member of the Oracle ASM administration group.

To use Oracle ACFS replication functionality, the disk group compatibility attributes for ASM and ADVM must be set to 12.2 or higher for the disk groups that contain the primary and standby file systems.

To use Oracle ACFS replication on Solaris Sparc hardware, the system must be running Solaris 10 update 8 or later.

To configure replication and manage replicated Oracle ACFS file systems, use the `acfsutil repl` command-line functions.

See Also:

- [Using Replication with Auditing, Encryption, and Security](#) (page 11-41) for information about using Oracle ACFS replication with other Oracle ACFS functionality
- [acfsutil size](#) (page 16-164) for information about resizing an Oracle ACFS file system
- [Oracle ACFS Tagging](#) (page 11-39) for information on tagging
- [About Privileges for Oracle ASM](#) (page 3-28) for information about user privileges for Oracle ASM
- [About Using Oracle ACFS Command-Line Tools](#) (page 16-2) for information about running Oracle ACFS `acfsutil` commands
- [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
- [Replicating Oracle ACFS File Systems](#) (page 16-17) and [Oracle ACFS Command-Line Tools for Replication](#) (page 16-63) for information about using the `acfsutil repl` command-line functions to configure replication and manage replicated Oracle ACFS file systems
- [Configuring Oracle ACFS Snapshot-Based Replication](#) (page 18-1) for information about converting an existing Oracle ACFS replication environment to the snapshot-based replication introduced in release 12.2.0.1
- For information about tuning your network, refer to the documentation at the MAA link on Oracle Technology Network:
 - <http://www.oracle.com/technetwork/database/features/availability/maa-096107.html>
 - Relevant information on tuning the network can be found in the *Data Guard Redo Transport & Network Configuration* paper

11.3.11 Oracle ACFS Tagging

Oracle ACFS tagging assigns a common naming attribute to a group of files.

Oracle ACFS Replication can use this tag to select files with a unique tag name for replication to a different remote cluster site. The tagging option avoids having to replicate an entire Oracle ACFS file system.

Oracle ACFS implements tagging with Extended Attributes. Some editing tools and backup utilities do not retain the Extended Attributes of the original file by default; you must set a specific switch. The following list describes the necessary requirements and switch settings for some common utilities to ensure Oracle ACFS tag names are preserved on the original file.

- The `cp` command requires flags to preserve tag names.

Install the `coreutils` library (version `coreutils-5.97-23.el5_4.1.src.rpm` or `coreutils-5.97-23.el5_4.2.x86_64.rpm` or later) on Linux to install versions of the `cp` command that supports Extended Attribute preservation with the `--preserve=xattr` switch and the `mv` command that supports Extended Attribute preservation without any switches.

`cp` does not preserve tag names assigned to symbolic link files.

The `cp` switches required to preserve tag names on files and directories are:

- Linux: `--preserve=xattr`
 - Solaris: `-@`
 - AIX: `-U`
 - Windows: no switch necessary
- The `cpio` file transfer utility requires flags to preserve tag names.

The `cpio` switches required to preserve tag names on files and directories are:

 - Linux: `cpio` does not preserve tag names
 - Solaris: `-@` is required to preserve or restore tag names for files and directories, but does not preserve tag names for symbolic link files
 - AIX: `-U` is required to preserve or restore tag names for files and directories, but does not preserve tag names for symbolic link files
 - Windows: not available
 - `emacs` requires that the `backup-by-copying` option is set to a non-nil value to preserve tag names on the original file name rather than a backup copy. This option must be added to the `.emacs` file.
 - The `pax` file transfer utility requires flags to preserve tag names.

The `pax` switches required to preserve tag names on files and directories are:

 - Linux: `pax` does not preserve tag names
 - Solaris: `-@` is required to preserve or restore tag names for files and directories, but does not preserve tag names for symbolic link files
 - AIX: `-U` is required to preserve or restore tag names for files and directories, but does not preserve tag names for symbolic link files
 - Windows: not available
 - The `rsync` file transfer utility requires flags to preserve tag names.

The `rsync` switches required to preserve tag names on files and directories are:

 - Linux: `-X -l` are required to preserve tag names for files and directories, but these switches do not preserve tag names for symbolic link files
 - Solaris: `rsync` does not preserve tag names
 - AIX: not available

- Windows: not available
- The `tar` backup utility can have flags set on the command line to preserve tag names on a file. However, `tar` does not retain the tag names assigned to symbolic link files.

The `tar` backup utility on Windows currently provides no support to retain tag names as no switch exists to save Extended Attributes.

The `tar` switches required to preserve tag names on files and directories are:

- Linux: `--xattrs`
- Solaris: `-@`
- AIX: `-U`
- Windows: `tar` does not preserve tag names
- The `vim` or `vi` editors require the `set bkc=yes` option in the `.vimrc` (Linux) or `_vimrc` (Windows) file to make a backup copy of a file and overwrite the original. This preserves tag names on the original file.

To use Oracle ACFS tagging functionality on Linux, the disk group compatibility attributes for ASM and ADVM must be set to 11.2.0.2 or higher. To use Oracle ACFS tagging functionality on Windows, the disk group compatibility attributes for ASM and ADVM must be set to 11.2.0.3 or higher. To use Oracle ACFS tagging functionality on Solaris or AIX, the disk group compatibility attributes for ASM and ADVM must be set to 12.1 or higher.

See Also:

- [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - [Tagging Oracle ACFS File Systems](#) (page 16-16) and [Oracle ACFS Command-Line Tools for Tagging](#) (page 16-60) for information about using the `acfsutil tag` command-line functions to configure tagging and manage tagged Oracle ACFS file systems
 - [Oracle ACFS Tagging Generic Application Programming Interface](#) (page 17-12) for information about Oracle ACFS tagging application programming interfaces (APIs)
-

11.3.12 Using Replication with Auditing, Encryption, and Security

Auditing, encryption, and realm-based security features can be enabled on an Oracle ACFS file system on which replication has been configured. The replicated standby file system is secured with the same auditing, security, or encryption policies as the primary file system. For this replicated environment, the primary and standby file systems must both be 12.1 or higher installations. For more information about Oracle ACFS replication, refer to "[Oracle ACFS Replication](#) (page 11-36)".

To ensure successful replication, the standby file system must be a generic file system without auditing, encryption, or security metadata on it. Oracle ACFS does not support using a standby file system that once had security or encryption and then had

security or encryption removed. Additional conditions that must be met for Oracle ACFS auditing, encryption, and security are listed in this section.

Note the following about Oracle ACFS audited file systems:

- Before replicating an audit-enabled file system or auditing a replicated file system, auditing must be initialized on the standby file system.
- Auditing policies present on the primary file system are replicated to the standby and any policy actions taken on the primary file system are enacted on the standby file system.
- Two sets of audit trails are present on the standby file system. Trails from primary file system are replicated to the standby file system as ordinary files. File system activity may generate events on the standby file system, which are recorded in the audit trail for the standby file system. Audit trail names help distinguish the two sets of trails because they contain both the host name and FSID.

Note the following about Oracle ACFS encrypted file systems:

- Encrypted files on the primary file system remain encrypted on the standby file system with the same key and encryption parameters (algorithm and key length).
- Encryption operations done on the primary file system are replayed on the standby file system - on, off, and rekey.
- Encryption may be enabled before or after a file system is replicated. In either case, an encryption wallet is transparently created on the standby file system if one does not exist because `acfsutil encr init` has not been run on the standby file system.
- A password-protected wallet is not supported on the standby file system. If a PKCS wallet already exists on a site that is to be used as a standby file system, the administrator must use the `acfsutil keystore migrate` command to transfer all keys to an SSO wallet.

Note the following about Oracle ACFS secured file systems:

- Standby file systems should be initialized for security before replicating a security enabled file system.
- The rules, rule sets and realms are replicated to the standby file system and same policies exist on the standby file system. In terms of the policies and protection of files, the standby file system is exactly same.
- Replication can be enabled on a security enabled file system or security can be enabled on a replicated file system. As part of security preparation, security is also enabled on the standby file system.
- Having security and replication together on a file system does not require any extra user intervention or additional steps.
- A different set of security administrators or security administrator groups can be set up on the standby file system.

11.3.13 Oracle ACFS Plugins

The Oracle ACFS plugin functionality enables a user space application to collect *just-in-time* Oracle ACFS file and Oracle ADVM volume metrics from the operating system environment.

Applications can use the Oracle ACFS plug-in infrastructure to create customized solutions that extend the general application file metric interfaces to include detailed Oracle ACFS file system and volume data.

The Oracle ACFS plug-in functionality can be enabled on separate Oracle ACFS file systems mounted on a standalone host or on one or more nodes of an Oracle Grid cluster where the Oracle ACFS file system is mounted. This functionality enables message communication between a node-local plugin enabled Oracle ACFS file system and an associated user space application module using Oracle ACFS plug-in application programming interfaces (APIs).

The plugin message APIs support both polling and posting message delivery models and multiple message payload types.

See Also:

- [Oracle ACFS Command-Line Utilities](#) (page 16-143) for information about Oracle ACFS plugin commands
 - [Oracle ACFS Plug-in Generic Application Programming Interface](#) (page 17-8) for information about the Oracle ACFS plug-in application programming interface
-
-

11.3.14 Oracle ACFS Accelerator Volume

Using an accelerator volume can improve performance by reducing the time to access and update Oracle ACFS metadata. You should create the accelerator volume on a disk group with storage that is significantly faster than the storage of the primary volume. For example, Solid State Disk (SSD) storage could be used. Oracle ADVM volumes are created with the ASMCMD `volcreate` command. For information about the `volcreate` command, refer to "[volcreate](#) (page 15-2)".

The recommended size of the accelerator volume depends on the workload. It is especially helpful for files with many extents, especially if that extent metadata is updated frequently. You can use the `acfsutil info file` command to view a report on a file's extents. Database files generally have many extents and when Oracle ACFS snapshots are in use, the extent metadata is updated frequently. A workload that greatly benefits from an accelerator is a compressed file system.

If Oracle ACFS cannot allocate space on the accelerator for critical metadata, then that metadata is stored on the primary volume instead. Depending on the frequency of metadata updates, it can have a disproportionate impact on performance. If the slow metadata is written in the same transaction as the fast metadata, then the slow metadata brings the performance of the entire operation down.

The recommended starting accelerator size is minimally 0.6% of the size of the file system. If many snapshots are in use representing several points in time for a database workload, the recommendation is an additional 0.4% per snapshot. For example, a file system with 5 snapshots may need an accelerator whose size is 2.6% of the size of the

primary volume. `acfsutil size` can be configured to automatically grow the accelerator as needed along with the primary volume. The accelerator increases in units of 64 mega bytes. The minimum size of the accelerator volume is 256 M. `mkfs` requires that the initial accelerator size be at least 0.4% of the size of the primary volume

The accelerator volume is linked to the primary volume specified with the `mkfs` command. When mounting a file system, only the primary volume is specified. If the accelerator volume becomes inaccessible for any reason after a file system with the volume is mounted, then the file system is taken offline. Only one storage accelerator volume can be associated with an Oracle ACFS file system. After an accelerator volume is associated with a file system, the volume cannot be disassociated from the file system.

The accelerator volume can be created on Linux environments with the `-a` option of the `mkfs` command. To use the `-a` option, the value of `COMPATIBLE.ADVM` must be at least 12.2. For information about the `mkfs` command, refer to "[mkfs](#) (page 16-28)".

11.3.15 Oracle ACFS NAS Maximum Availability eXtensions

Oracle ACFS NAS Maximum Availability eXtensions (Oracle ACFS NAS MAX) is a set of extensions that provide High Availability Extensions for Common NAS Protocols, such as NFS and SMB.

When using these extensions, the protocol in question is running in high availability mode, enabling the protocol to move between nodes in an Oracle RAC cluster. This functionality provides a way to address a single point of failure for a given protocol, so that if at least one node of the cluster is available, then the protocol is available. In addition to providing for high availability, the extensions provide for integration with common NAS protocols and the Oracle ACFS stack, enabling administrators to easily utilize these protocols without creating additional infrastructure. Note that the Oracle ACFS NAS Maximum Availability eXtensions functionality adds value to existing OS NAS protocol implementations, but does not replace them.

Oracle ACFS High Availability Network File System

High Availability Network File System (HANFS) for Oracle Grid Infrastructure provides uninterrupted service of NFS v2, v3, or v4 exported paths by exposing NFS exports on Highly Available Virtual IPs (HAVIP) and using Oracle Clusterware agents to ensure that the VIPs and NFS exports are always online. While base NFS supports file locking, HANFS does not support NFS file locking.

Note:

- This functionality relies on a working NFS server configuration available on the host computer. You must configure the NFS server before attempting to use the Oracle ACFS NFS export functionality.
 - This functionality is not available on Windows.
 - This functionality is not supported in Oracle Restart configurations.
 - The HAVIP cannot be started until at least one file system export resource has been created for it.
-
-

To set up High Availability NFS for Oracle Grid Infrastructure, perform the following steps:

1. Add and register a new HAVIP resource.

For example:

```
# srvctl add havip -id hreexports -address my_havip_name
```

In the example, *my_havip_name* is mapped in the domain name server (DNS) to the VIP address and is used by the client systems when mounting the file system.

The initial processing of `srvctl add havip` ensures that:

- The address being used is static, not dynamic
- Any DNS names resolve to only one host, not round-robin multiple DNS resolutions
- The network resource and provided IP address and resolved name are in the same subnet
- The name is not in use

SRVCTL creates the appropriate HAVIP name using the *id*, ensuring it is unique. As a final validation step, SRVCTL ensures that the *network* resource (if provided) of `ora.net#.network` exists. After this step, SRVCTL adds a new `havip` of type `ora.havip.type` with the name of `ora.id.havip`. In this example, the name is `ora.hreexports.havip`.

Next SRVCTL modifies HAVIP start dependencies, such as *active dispersion*; sets the stop dependencies; and ensures the *description* attribute (if provided) is appropriately set.

2. Create a shared Oracle ACFS file system.

High Availability NFS for Oracle Grid Infrastructure operates only with Oracle ACFS file systems configured for clusterwide accessibility and does not support Oracle ACFS file systems configured for access on particular subsets of cluster nodes. High Availability NFS is not supported with non-Oracle ACFS file systems.

3. Register the Oracle ACFS file system.

For example:

```
$ srvctl add filesystem -device /dev/asm/dlvolume1-295 -volume VOLUME1 \
  -diskgroup HR_DATA -mountpath /oracle/cluster1/acfs1
```

4. Create an Oracle ACFS file system export resource.

For example:

```
# srvctl add exportfs -id hreexports -path /oracle/cluster1/acfs1 -name hreexport1
```

After the file system export resource has been created, then you can start the HAVIP created in step 1 to export the file system using the `srvctl start havip` command.

The NFS mount option `FSID` is added to any export options, utilizing the `FSID` of the underlying Oracle ACFS file system plus a unique identifier. This `FSID` option provides for reliable fail over between nodes and allows the usage of snapshot mounting.

The default mount and export options for configured exports are the defaults for the NFS server.

Relative paths that are fully-qualified are converted to absolute paths. Relative paths that are not fully-qualified are not accepted as an export path.

VIPs attempts to find the best server to run on based on available file systems and other running VIPs, but this dispersion only occurs during CSS membership change events, such as a node joining or leaving the cluster.

Note:

It is not recommended to start and stop exports individually; this functionality should be provided through the start and stop operations of HAVIP.

When HAVIP is not running, exports can exist on different nodes. After the associated HAVIP is started, the exports gather on a single node.

Clients that are using an export that is stopped while HAVIP is running raise the NFS error `estale`, and must dismount and remount the file system.

When mounting an HANFS exported file system on a client, the following `CLIENT` mount options are recommended:

```
hard,intr,retrans=10000
```

Oracle ACFS HANFS with NFS Locks

Oracle ACFS HANFS now supports HANFS NFS v4 with NFS Locks. This functionality is only available on specific operating system (OS) platforms. To activate this functionality, additional steps must be performed after the Oracle Grid Infrastructure software is installed. Note that after these steps are completed, the OS NFS server functionality of the cluster is managed by the Oracle Clusterware stack. In addition, the location of certain OS NFS configuration files will be moved from their default location to a designated Oracle ACFS file system.

Some common tasks are:

- Activate: `acfshanfs addnode`
- Uninstall: `acfshanfs uninstall`
- Check the installation status: `acfshanfs installed`
- Check if this platform is supported: `acfshanfs supported`

When activating the HANFS v4 lock functionality, the following command must be run on each node:

```
# grid_home/bin/acfshanfs addnode -nfsv4lock -volume volume_device
```

The volume is formatted with an Oracle ACFS file system and mounted on a designated Oracle ACFS clusterware mount point. For example on Linux:

```
/dev/asm/nfs-81 on /var/lib/nfs type acfs (rw)
```

Restrictions on the Oracle ADVM volume include:

- No previously existing Oracle ACFS resource should exist for this new Oracle ADVM volume.
- No Oracle ACFS file system should exist on this Oracle ADVM volume.
- This Oracle ADVM volume should not be in use anywhere in the cluster.

When Oracle HANFS v4 lock functionality is activated, there are differences from normal HANFS operations. The differences are noted in the following list:

- The OS NFS server is under Oracle Clusterware control through the `ora.netstorageservice` resource. When starting and stopping the Oracle Clusterware stack, the OS NFS server is also started and stopped.
- This resource has a dependency on an Oracle ACFS file system:
`ora.data_hostname.nfs.acfs`
The *hostname* is the hostname of the first node on which the setup for Oracle HANFS locking has been run.
- Only Oracle HANFS should be used to export NFS file systems from the Oracle RAC cluster. The NFS server is configured and moved around the Oracle RAC cluster; only file systems exported by Oracle HANFS are accessible when the NFS server has migrated to an alternate cluster node.
- When locking is initialized, Oracle HANFS exports are run from only a single node, unlike non-locking mode, where Oracle HANFS exports are distributed throughout the cluster.
- On client nodes, mount the file system specifying NFS v4 as the NFS version. This prevents the server from defaulting to NFS v3, and enables support for the NFS v4 locking functionality.

After High Availability Locking is activated, control of HANFS with locking is the same as described previously in this section.

Oracle ACFS HANFS with High Availability SMB

Oracle ACFS supports High Availability Samba (SMB), also known as CIFS (Common Internet File System) in previous Microsoft implementations. This protocol is commonly used to interface with Microsoft servers and Active Directory Domains and is supported by various operating system (OS) implementations. However, Oracle ACFS High Availability SMB requires the Microsoft SMB implementation or Samba.

Note the following:

- Samba is available from www.samba.org
- Ensure that Samba or SMB is correctly configured on your host OS before attempting to utilize High Availability SMB.
- High Availability SMB is not supported in Oracle Restart mode.
- After adding an HAVIP resource, an SMB Export resource must also be added; otherwise, the HAVIP resource does not start.
- For highest performance and best results, ensure that both server and client are using SMB3. Note the following:
 - Use the newest version of Samba, v4 or later.
 - Use the latest Microsoft OS version (2012 or later). To check the SMB version, use the Powershell cmdlet `Get-SmbConnection` command.
 - Previous versions of SMB require that the client must remount the SMB export after a storage failure.

- Similar to HANFS, options may be specified on the command line and are passed to the host operating system. Appropriate error messages are passed back. If no options are provided to the SRVCTL command, the following default options apply:
 - Windows: READ Access for Everyone
 - Linux, Solaris, and AIX: Read Only, Browsable = True
- Supported Option Sets:
 - Windows: Any options supported by the `net . exe` command.
 - Linux, Solaris, or AIX: Any options supported by the Samba configuration stanza.

To set up High Availability SMB for Oracle Grid Infrastructure, perform the following steps:

1. Add and register a new HAVIP resource.

For example:

```
# srvctl add havip -id hreexports -address my_havip_name
```

In the example, `my_havip_name` is mapped in the domain name server (DNS) to the VIP address and is used by the client systems when mounting the file system.

The initial processing of `srvctl add havip` ensures that:

- The address being used is static, not dynamic
- Any DNS names resolve to only one host, not round-robin multiple DNS resolutions
- The network resource and provided IP address and resolved name are in the same subnet
- The name is not in use

SRVCTL creates the appropriate HAVIP name using the `id`, ensuring it is unique. As a final validation step, SRVCTL ensures that the `network` resource (if provided) of `ora.net#.network` exists. After this step, SRVCTL adds a new `havip` of type `ora.havip.type` with the name of `ora.id.havip`. In this example, the name is `ora.hreexports.havip`.

Next SRVCTL modifies HAVIP start dependencies, such as `active dispersion`; sets the stop dependencies; and ensures the `description` attribute (if provided) is appropriately set.

2. Create a shared Oracle ACFS file system.

High Availability SMB for Oracle Grid Infrastructure operates only with Oracle ACFS file systems configured for clusterwide accessibility and does not support Oracle ACFS file systems configured for access on particular subsets of cluster nodes. High Availability NFS is not supported with non-Oracle ACFS file systems.

3. Register the Oracle ACFS file system.

For example:

```
$ srvctl add filesystem -device /dev/asm/dlvolume1-295 -volume VOLUME1 \
-diskgroup HR_DATA -mountpath /oracle/cluster1/acfs1
```

4. Create an Oracle ACFS file system export resource.

For example:

```
# srvctl add exportfs -id hrexperts -path /oracle/cluster1/acfs1 -name hrexpert1
-type SMB
```

After the file system export resource has been created, then you can start the HAVIP created in step 1 to export the file system using the `srvctl start havip` command.

During the start of the resource, the Oracle ACFS Export resource creates a Samba configuration file (Linux, Solaris, or AIX) or runs the `net . exe` binary to export the file system.

VIPs attempts to find the best server to run on based on available file systems and other running VIPs, but this operation only occurs during CSS membership change events, such as a node joining or leaving the cluster.

Note:

- It is not recommended to start and stop exports individually; this functionality should be provided through the start and stop operations of HAVIP.
 - When HAVIP is not running, exports can exist on different nodes. After the associated HAVIP is started, the exports gather on a single node.
-
-

See Also:

- [Creating an Oracle ACFS File System](#) (page 16-5) for information about creating an Oracle ACFS file system
 - *Oracle Clusterware Administration and Deployment Guide* for information about the `srvctl add filesystem` command
 - *Oracle Clusterware Administration and Deployment Guide* for information about Oracle Clusterware resources
 - *Oracle Real Application Clusters Administration and Deployment Guide* for more information about SRVCTL commands
-
-

11.4 Overview of Oracle ASM Dynamic Volume Manager

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 send 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 failures.

Oracle ADVM extends Oracle ASM by providing a disk driver interface to Oracle ASM storage allocated as Oracle ADVM volume files. You can use Oracle ADVM to create virtual disks that contain file systems. These file systems contained on Oracle ADVM volumes are able 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) communicates with 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'.

Note:

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.

See Also:

- [Managing Oracle ADVM Volumes in a Disk Group](#) (page 4-21) for information about the `ALTER DISKGROUP ADD | RESIZE | DROP | ENABLE | DISABLE | MODIFY VOLUME` SQL statements that manage Oracle ADVM volumes
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about ASMCMD command-line tools for managing Oracle ADVM volumes
 - [Using Views to Display Oracle ACFS Information](#) (page 12-1) for information about `V$ASM` views to display information about Oracle ADVM volumes
 - [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1) for information about Oracle ACFS operating system utilities for managing file systems and volume device files
 - *Oracle Database SQL Language Reference* for information about the `ALTER DISKGROUP` SQL statement
-
-

Using Views to Display Oracle ACFS Information

Dynamic views display important information about Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM).

This chapter contains information about using dynamic views to display Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) information.

- [Views Containing Oracle ACFS Information](#) (page 12-1)
- [Oracle ACFS Support for Oracle Database File Mapping Views](#) (page 12-6)

See Also:

Oracle Database Reference for information about all of the V\$ASM* dynamic performance views

12.1 Views Containing Oracle ACFS Information

You can use the views in [Table 12-1](#) (page 12-1) to obtain information about Oracle Automatic Storage Management Cluster File System (Oracle ACFS). These views are accessible from the Oracle ASM instance.

Note:

- To display information about Oracle ACFS file systems or volumes that are located on nodes in an Oracle Flex ASM configuration, you must connect to the Oracle ASM proxy instance instead of the local Oracle ASM instance.
 - When viewing space usage values in Oracle ACFS views on Windows systems, the values may differ from sizes in Windows folders. The mechanism used by **Folder Properties** on Windows systems only accounts for files and should be considered an approximate value.
-

Table 12-1 Oracle ASM dynamic views for Oracle ACFS information

View	Description
V\$ASM_ACFS_ENCRYPTION_INFO	Contains encryption information for each Oracle ACFS file system.

Table 12-1 (Cont.) Oracle ASM dynamic views for Oracle ACFS information

View	Description
V\$ASM_ACFS_SECURITY_INFO	Contains security realm information for each Oracle ACFS file system.
V\$ASM_ACFS_SEC_ADMIN	Contains information about Oracle ACFS security administrator in the cluster.
V\$ASM_ACFS_SEC_CMDRULE	Contains information about Oracle ACFS security command rules for each Oracle ACFS file system.
V\$ASM_ACFS_SEC_REALM	Contains information about every realm in the Oracle ACFS security realm for each Oracle ACFS file system.
V\$ASM_ACFS_SEC_REALM_FILTER	Contains information about every filter in the Oracle ACFS security realm for each Oracle ACFS file system. A filter is defined as <code>commandrule/ruleset</code> pair in a realm.
V\$ASM_ACFS_SEC_REALM_GROUP	Contains information about every group in the Oracle ACFS security realm for each Oracle ACFS file system.
V\$ASM_ACFS_SEC_REALM_USER	Contains information about every user in the Oracle ACFS security realm for each Oracle ACFS file system.
V\$ASM_ACFS_SEC_RULE	Contains information about every Oracle ACFS security rule for each Oracle ACFS file system.
V\$ASM_ACFS_SEC_RULESET	Contains information about every Oracle ACFS security ruleset for each Oracle ACFS file system.
V\$ASM_ACFS_SEC_RULESET_RULE	Contains information about every rule in Oracle ACFS security ruleset for each Oracle ACFS file system.
V\$ASM_ACFSREPL	Contains information for Oracle ACFS file systems that are initialized for replication. This view only contains records for Oracle ASM releases prior to 12.2. To display Oracle ACFS replication information for Oracle ASM releases 12.2 or higher, use the <code>acfsutil repl info</code> command.
V\$ASM_ACFSREPLTAG	Contains replicated tag information for Oracle ACFS file systems that are initialized for replication. This view only contains records for Oracle ASM releases prior to 12.2. To display Oracle ACFS replication information for Oracle ASM releases 12.2 or higher, use the <code>acfsutil repl info</code> command.
V\$ASM_ACFSSNAPSHOTS	Contains snapshot information about every mounted Oracle ACFS file system.
V\$ASM_ACFSTAG	Contains all tag name information about files on all mounted Oracle ACFS file systems
V\$ASM_ACFSVOLUMES	Contains information about mounted Oracle ADVM volumes, correlated with V\$ASM_FILESYSTEM.
V\$ASM_FILESYSTEM	Contains information about 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.

Table 12-1 (Cont.) Oracle ASM dynamic views for Oracle ACFS information

View	Description
V\$ASM_VOLUME_STAT	Contains information about statistics for each Oracle ADVM volume.

Example 12-1 Viewing encryption information in V\$ASM_ACFS_ENCRYPTION_INFO

This example shows information displayed from the V\$ASM_ACFS_ENCRYPTION_INFO view. The FS_NAME column contains the mount point. The VOL_DEVICE contains the name of the Oracle ADVM device.

```
SELECT SUBSTR(fs_name,1,24) FILESYSTEM, SUBSTR(vol_device,1,20) DEVICE, set_status,
       enabled_status, algorithm, key_length, last_rekey_time FROM V$ASM_ACFS_ENCRYPTION_INFO;
```

FILESYSTEM	DEVICE	SET_STA	ENABLED_	ALGORIT	KEY_LEN	LAST_REKE
/acfsmounts/acfs1	/dev/asm/volume1-228	YES	ENABLED	AES	192	

Example 12-2 Viewing security information in V\$ASM_ACFS_SECURITY_INFO

This example shows information displayed from the V\$ASM_ACFS_SECURITY_INFO view. The FS_NAME column contains the mount point. The VOL_DEVICE contains the name of the Oracle ADVM device.

```
SELECT SUBSTR(fs_name,1,24) FILESYSTEM, SUBSTR(vol_device,1,20) DEVICE, prepared_status,
       enabled_status FROM V$ASM_ACFS_SECURITY_INFO;
```

FILESYSTEM	DEVICE	PREPARE	ENABLED_
/acfsmounts/acfs1	/dev/asm/volume1-228	YES	ENABLED

Example 12-3 Viewing security rules information in V\$ASM_ACFS_SEC_RULE

This example shows information displayed from the V\$ASM_ACFS_SEC_RULE view. The FS_NAME column contains the mount point.

```
SELECT SUBSTR(rule_name,1,24) rule, SUBSTR(rule_type,1,12) type,
       SUBSTR(rule_value,1,16) value, SUBSTR(fs_name,1,24) filesystem
FROM V$ASM_ACFS_SEC_RULE;
```

RULE	TYPE	VALUE	FILESYSTEM
medHistRule1a	TIME	22:00:00	/acfsmounts/acfs1
medHistRule1c	TIME	08:00:00	/acfsmounts/acfs1
medHistRule1b	USERNAME	medMaintenance	/acfsmounts/acfs1
medHistRule1d	USERNAME	medBrowse	/acfsmounts/acfs1
SYSTEM_RULE_Auditor	GROUPNAME	myaudit_mgr_grp	/acfsmounts/acfs1
SYSTEM_RULE_AuditManager	GROUPNAME	myauditor_grp	/acfsmounts/acfs1
SYSTEM_RULE_Always	TIME	00:00:00	/acfsmounts/acfs1

Example 12-4 Viewing security ruleset information in V\$ASM_ACFS_SEC_RULESET

This example shows information displayed from the V\$ASM_ACFS_SEC_RULESET view. The FS_NAME column contains the mount point.

```
SELECT SUBSTR(ruleset_name,1,36) ruleset, ruleset_option r_option, SUBSTR(fs_name,1,24) filesystem
FROM V$ASM_ACFS_SEC_RULESET;
```

RULESET	R_OPTION	FILESYSTEM

```

medRuleSet1          ALL_TRUE /acfsmounts/acfs1
medRuleSet2          ALL_TRUE /acfsmounts/acfs1
SYSTEM_RULESET_Auditor  ALL_TRUE /acfsmounts/acfs1
SYSTEM_RULESET_AuditManager ALL_TRUE /acfsmounts/acfs1
SYSTEM_RULESET_AuditMgr_Auditor ANY_TRUE /acfsmounts/acfs1
SYSTEM_RULESET_AlwaysDeny ANY_TRUE /acfsmounts/acfs1

```

Example 12-5 Viewing security ruleset information in V\$ASM_ACFS_SEC_RULESET_RULE

This example shows information displayed from the V\$ASM_ACFS_SEC_RULESET_RULE view.

```

SELECT SUBSTR(ruleset_name,1,36) ruleset, substr(rule_name,1,24) rule,
       SUBSTR(fs_name,1,36) filesystem FROM V$ASM_ACFS_SEC_RULESET_RULE;

```

RULESET	RULE	FILESYSTEM
medRuleSet1	medHistRule1a	/acfsmounts/acfs1
medRuleSet1	medHistRule1b	/acfsmounts/acfs1
medRuleSet2	medHistRule1c	/acfsmounts/acfs1
medRuleSet2	medHistRule1d	/acfsmounts/acfs1
SYSTEM_RULESET_Auditor	SYSTEM_RULE_Auditor	/acfsmounts/acfs1
SYSTEM_RULESET_AuditManager	SYSTEM_RULE_AuditManager	/acfsmounts/acfs1
SYSTEM_RULESET_AuditMgr_Auditor	SYSTEM_RULE_Auditor	/acfsmounts/acfs1
SYSTEM_RULESET_AuditMgr_Auditor	SYSTEM_RULE_AuditManager	/acfsmounts/acfs1
SYSTEM_RULESET_AlwaysDeny	SYSTEM_RULE_Always	/acfsmounts/acfs1

Example 12-6 Viewing snapshot information in V\$ASM_ACFSSNAPSHOTS

This example 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.

```

SELECT SUBSTR(FS_NAME,1,24) FILESYSTEM, SUBSTR(VOL_DEVICE,1,22) DEVICE,
       SUBSTR(SNAP_NAME,1,12) SNAPSHOT, CREATE_TIME TIME, SUBSTR(PARENT,1,10) PARENT,
       SUBSTR(TYPE,1,4) TYPE FROM V$ASM_ACFSSNAPSHOTS;

```

FILESYSTEM	DEVICE	SNAPSHOT	TIME	PARENT	TY
/acfsmounts/acfs1	/dev/asm/volume1-229	snapr	13-MAR-12	NULL	RW
/acfsmounts/acfs1	/dev/asm/volume1-229	snapr_child	13-MAR-12	snapr	RW
/acfsmounts/acfs2	/dev/asm/volume2-321	snapro	13-MAR-12	NULL	RO

Example 12-7 Viewing volume information with V\$ASM_ACFSVOLUMES

This example shows information displayed from the V\$ASM_ACFSVOLUMES view. The PRIMARY_VOL column contains TRUE if the volume is the primary volume for the file system.

```

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
/acfsmounts/acfs1	/dev/asm/volume1-228	TRUE	1024000	578626.522
/acfsmounts/acfs2	/dev/asm/volume2-375	TRUE	1024000	685761.463
...				

Example 12-8 Viewing volume information with V\$ASM_FILESYSTEM

This example 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. With 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 `acfschkdsk` command run on it.

```
SELECT fs_name, available_time, block_size, state, corrupt FROM V$ASM_FILESYSTEM;
```

FS_NAME	AVAILABLE	BLOCK_SIZE	STATE	CORRUPT
/acfsmounts/acfs1	19-JUL-09	4	AVAILABLE	FALSE
/acfsmounts/acfs2	19-JUL-09	4	AVAILABLE	FALSE

Example 12-9 Viewing volume information with V\$ASM_VOLUME

This example shows information displayed from the V\$ASM_VOLUME view for volumes contained in the DATA disk group.

```
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	/acfsmounts/acfs1
DATA	VOLUME2	/dev/asm/volume2-375	/acfsmounts/acfs2

Example 12-10 Viewing volume information with V\$ASM_VOLUME_STAT

This example 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.

```
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

Example 12-11 Viewing tag name information with V\$ASM_ACFSTAG

This example shows tag names for the /acfsmounts/acfs1 file system displayed from the V\$ASM_ACFSTAG view.

```
SELECT SUBSTR(TAG_NAME,1,8) TAG_NAME, SUBSTR(FS_NAME,1,20) FS_NAME,
SUBSTR(PATH_NAME,1,42) PATH_NAME FROM V$ASM_ACFSTAG WHERE TAG_NAME='tag5';
```

TAG_NAME	FS_NAME	PATH_NAME
tag5	/acfsmounts/acfs1	/acfsmounts/acfs1/d1/d2/d3/d4/d5/f6
tag5	/acfsmounts/acfs1	/acfsmounts/acfs1/d1/d2/d3/d4/d5
tag5	/acfsmounts/acfs1	/acfsmounts/acfs1/d1/d2/d3/d4
tag5	/acfsmounts/acfs1	/acfsmounts/acfs1/d1/d2/d3
tag5	/acfsmounts/acfs1	/acfsmounts/acfs1/d1/d2
tag5	/acfsmounts/acfs1	/acfsmounts/acfs1/d1

```
tag5      /acfsmounts/acfs1      /acfsmounts/acfs1/fl
tag5      /acfsmounts/acfs1      /acfsmounts/acfs1
```

See Also:

- *Oracle Database Reference* for information about the V\$ASM_ACFS_ENCRYPTION_INFO view
 - *Oracle Database Reference* for information about the V\$ASM_ACFS_SECURITY_INFO view
 - *Oracle Database Reference* for information about the V\$ASM_ACFSSNAPSHOTS view
 - *Oracle Database Reference* for information about the V\$ASM_ACFSVOLUMES view
 - *Oracle Database Reference* for information about the V\$ASM_FILESYSTEM view
 - *Oracle Database Reference* for information about the V\$ASM_VOLUME view
 - *Oracle Database Reference* for information about the V\$ASM_VOLUME_STAT view
 - *Oracle Database Reference* for information about the V\$ASM_ACFSTAG view
-
-

12.2 Oracle ACFS Support for Oracle Database File Mapping Views

Oracle ACFS supports Oracle Database file mapping views to the Oracle ASM device level.

Note:

This feature is available starting with Oracle Database 12c Release 1 (12.1.0.2).

The following database mapping views are supported by Oracle ACFS:

- V\$MAP_FILE
- V\$MAP_FILE_EXTENT
- V\$MAP_ELEMENT
- V\$MAP_FILE_IO_STACK

These V\$MAP views are only refreshed by executing the procedure DBMS_STORAGE_MAP.MAP_ALL. The Oracle ACFS file mapping interface does not utilize the external fmpu1 process or its supporting libraries.

Note:

- Oracle ACFS does not provide support for Oracle Database file mapping on Windows.
- Oracle ACFS does not provide support for the V\$MAP_SUBELEMENT view.

Before running any queries on the V\$MAP views, ensure that the FILE_MAPPING initialization is set to TRUE, then run the DBMS_STORAGE_MAP.MAP_ALL procedure to build the mapping information for the entire I/O subsystem associated with the database. For example, connect as SYSDBA to the database instance and run the following:

```
SQL> ALTER SYSTEM SET file_mapping=true;

SQL> EXEC DBMS_STORAGE_MAP.MAP_ALL(10000);
```

The SQL statements in [Example 12-12](#) (page 12-7) to [Example 12-15](#) (page 12-8) are run from the Oracle Database instance.

Example 12-12 Viewing Oracle ASM information with V\$MAP_ELEMENT

This example displays information from the V\$MAP_ELEMENT view.

```
SQL> SELECT ELEM_NAME, ELEM_IDX, ELEM_TYPE, ELEM_SIZE, ELEM_DESCR
       FROM V$MAP_ELEMENT;
```

ELEM_NAME	ELEM_IDX	ELEM_TYPE	ELEM_SIZE	ELEM_DESCR
+/dev/xvdd1	0	ASMDISK	117184512	TEST_0001
+/dev/xvdc1	1	ASMDISK	117184512	TEST_0000

Example 12-13 Viewing Oracle ACFS Data File Information with V\$MAP_FILE

This example displays information from the V\$MAP_FILE view.

```
SQL> SELECT FILE_NAME, FILE_MAP_IDX, FILE_TYPE, FILE_STRUCTURE, FILE_SIZE,
       FILE_NEXTS FROM V$MAP_FILE WHERE REGEXP_LIKE(FILE_NAME, '*users01.dbf');
```

FILE_NAME	FILE_MAP_IDX	FILE_TYPE	FILE_STRU	FILE_SIZE	FILE_NEXTS
/dbdata1/orcl/users01.dbf	4	DATAFILE	FILE	10256	41

Example 12-14 Viewing Element and File Offset Information with V\$MAP_FILE_EXTENT

This example displays the element offset versus file offset information for each extent with V\$MAP_FILE_EXTENT, specifying FILE_MAP_IDX equal to 4, which is the file map index of the /dbdata/orcl/users01.dbf file.

```
SQL> SELECT FILE_MAP_IDX, EXT_NUM, EXT_ELEM_OFF, EXT_SIZE, EXT_FILE_OFF,
       EXT_TYPE, ELEM_IDX FROM V$MAP_FILE_EXTENT WHERE FILE_MAP_IDX=4;
```

FILE_MAP_IDX	EXT_NUM	EXT_ELEM_OFF	EXT_SIZE	EXT_FILE_OFF	EXT_TY	ELEM_IDX
4	0	58105664	192	0	DATA	0
4	1	58154752	256	192	DATA	1
4	2	58089472	256	448	DATA	0

...

```

          4          39      58140928      256          9920 DATA          1
          4          40      58108160       88          10176 DATA          0
41 rows selected.

```

Example 12-15 Viewing Extent Information With V\$MAP_FILE_IO_STACK

This example displays information from V\$MAP_FILE_IO_STACK specifying FILE_MAP_IDX equal to 4. The V\$MAP_FILE_IO_STACK view is similar to V\$MAP_FILE_EXTENT, but the display groups contiguous extents which are on the same device or element and of the same size.

```
SQL> SELECT FILE_MAP_IDX, ELEM_IDX, CU_SIZE, STRIDE, NUM_CU, ELEM_OFFSET,
          FILE_OFFSET FROM V$MAP_FILE_IO_STACK WHERE FILE_MAP_IDX=4;
```

FILE_MAP_IDX	ELEM_IDX	CU_SIZE	STRIDE	NUM_CU	ELEM_OFFSET	FILE_OFFSET
4	0	256	1024	10	58089472	448
4	0	192	0	1	58105664	0
4	0	256	1024	9	58105856	960
4	0	88	0	1	58108160	10176
4	1	256	1024	10	58138624	704
4	1	256	1024	10	58154752	192

6 rows selected.

See Also:

- [Oracle Database File Mapping for Oracle ASM Files](#) (page 5-11) for information about file mapping to Oracle ASM files
 - *Oracle Database Administrator's Guide* for information about the Oracle Database File Mapping Interface
 - *Oracle Database Reference* for details about the V\$MAP views
-

Administering Oracle ACFS with Oracle Enterprise Manager

Oracle Enterprise Manager Cloud Control provides tools for administering Oracle Automatic Storage Management Cluster File System (Oracle ACFS).

This chapter describes how to administer Oracle Automatic Storage Management Cluster File System (Oracle ACFS) with Oracle Enterprise Manager Cloud Control.

All Oracle ASM administration tasks begin with the Oracle Automatic Storage Management home page in Oracle Enterprise Manager Cloud Control.

Note:

To manage or monitor Oracle ACFS file systems or volumes that are located on nodes in an Oracle Flex ASM configuration, you must connect to the Oracle ASM proxy instance instead of the local Oracle ASM instance.

This chapter contains the following topics:

- [Managing Oracle ACFS File Systems and Volumes with Oracle Enterprise Manager](#) (page 13-2)
- [Managing Oracle ACFS Snapshots with Oracle Enterprise Manager](#) (page 13-4)
- [Managing Security Features with Oracle Enterprise Manager](#) (page 13-6)
- [Managing Encryption Features with Oracle Enterprise Manager](#) (page 13-16)
- [Managing Tagging Features with Oracle Enterprise Manager](#) (page 13-17)

See Also:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for information about the Oracle ASM home page and how to access that page
 - [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1) for information about using command-line tools to administer Oracle ACFS file systems
 - [Basic Steps to Manage Oracle ACFS Systems](#) (page 16-2) for information about the basic steps for creating Oracle ACFS file systems
 - [Managing Oracle Flex ASM](#) (page 19-1) for information about Oracle Flex ASM
-

13.1 Managing Oracle ACFS File Systems and Volumes with Oracle Enterprise Manager

This section discusses how to manage Oracle ACFS file systems and volumes systems with Oracle Enterprise Manager Cloud Control.

This section contains the following topics:

- [Creating Oracle ACFS File Systems and Volumes](#) (page 13-2)
- [Viewing and Modifying Oracle ACFS Volumes and File Systems](#) (page 13-3)

For more information about Oracle ACFS file systems and volumes, refer to "[About Oracle ACFS](#) (page 11-5)".

See Also:

- *Oracle Enterprise Manager Cloud Administration Guide*
 - *Oracle Enterprise Manager Administration*
-
-

13.1.1 Creating Oracle ACFS File Systems and Volumes

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.
3. Click **Create** in the **Volumes** tab of the **Disk Group** page.

The **Create ASM Volume** page displays.

Enter the volume name, 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.

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.
2. Select an existing volume in the disk group and click **Create ASM Cluster File System**.

The compatibility parameters `COMPATIBLE.ASM` and `COMPATIBLE.ADVM` must be set to 11.2 or higher for the disk group..

3. On the **Create ASM Cluster File System** page, enter the information to create a file system. You can optionally register and mount the file system.

You must 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 run the command when prompted. To register or mount a file system, you need root or administrator privileges.

After a directory has been selected, click **Show Command** to have Oracle Enterprise Manager generate the commands that you can run 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 Also:

- [Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2) for information about the Oracle ASM home page and how to access that page
 - [Mirroring, Redundancy, and Failure Group Options](#) (page 4-38) for information about redundancy settings
 - [Intelligent Data Placement](#) (page 4-26) for information about hot and cold disk regions
 - [ASMCMD Volume Management Commands](#) (page 15-1) for information about creating a volume with the ASMCMD `volcreate` command
 - [Disk Group Compatibility Attributes](#) (page 4-51) for information about disk group compatibility
 - [About the Oracle ACFS Mount Registry](#) (page 11-11) for information about the registry process
-
-

13.1.2 Viewing and Modifying Oracle ACFS Volumes and File Systems

To view and modify information about Oracle ACFS file systems, click the **ASM Cluster File System** link in the Oracle ASM home page.

The **ASM Cluster File System** tab in Oracle Enterprise Manager 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 the illustration. The mount point, snapshots, volume, and disk group are provided as links to further information.

To view information about a specific volume, click the volume name in the **Volume** column on the **ASM Cluster File System** page to display the **General** tab of the **ASM Volumes** page.

To view information about a file system, click the link in the **Mount Point** column on the **ASM Cluster File System** page.

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.

To view information about volumes in a disk group, click the **Volumes** tab at the top of the **Disk Group** page.

13.2 Managing Oracle ACFS Snapshots with Oracle Enterprise Manager

This section describes how to manage snapshots with Oracle Enterprise Manager.

- [Creating, Modifying, and Viewing Snapshots](#) (page 13-4)
- [Converting Snapshots Between Read-Only and Read-Write](#) (page 13-5)
- [Creating Child Snapshots from Existing Snapshots](#) (page 13-5)

For more information about Oracle ACFS snapshots, refer to "[About Oracle ACFS Snapshots](#) (page 11-11)".

13.2.1 Creating, Modifying, and Viewing Snapshots

To view and modify information about existing snapshots or create snapshots for a file system, perform the following steps.

1. Click the **ASM Cluster File System** tab on the Oracle ASM home page
2. Click a mount point link in the **Mount Point** column.
3. Click the **Snapshots** tab.

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**.

To create a snapshot, perform the following steps.

1. Click **Create** in the **Snapshots** page.
2. Complete the information on the **Create Snapshot** 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 to manually enter the generated operating system command.

To drill down in a snapshot directory, click the snapshot name in the **Snapshots** page to display the **Search and List Directory** page.

13.2.2 Converting Snapshots Between Read-Only and Read-Write

You can use edit snapshot to convert the snapshot attribute between Read Only and Read Write. To edit a snapshot, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

2. Click the file or directory in which you want to edit a snapshot.

Enterprise Manager displays ASM Cluster File System page.

3. Click the **Snapshots** tab.

4. Select the snapshot you want to edit, then click **Edit**.

The Snapshot page displays.

5. Change the Attribute. You can choose **Read Only** or **Read Write**.

6. Click **OK**.

13.2.3 Creating Child Snapshots from Existing Snapshots

Use the Create Child snapshot operation to create a point-in-time copy of an existing Oracle ACFS snapshot. Future changes to the parent snapshot are not inherited by the child snapshot. To create a child snapshot of an existing snapshot, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

2. Click the file or directory Mount Point in which you want to create a child snapshot.

Enterprise Manager displays ASM Cluster File System page.

3. Click the **Snapshots** tab.

4. Select the parent snapshot you want to use to create a child snapshot, then click **Create Child**.

The Create Child Snapshot on ASM Cluster File System page displays.

5. The Snapshot Name field is automatically populated with a default child snapshot name. You can replace the name with a name you designate.

6. Set the Attribute to **Read Only** or **Read Write**.

7. Optionally you can turn on the option to delete the oldest snapshot by turning on the delete option. Oracle Enterprise Manager tabulates the number of remaining snapshots before the maximum limit is reached.

8. Click **OK** to create the child snapshot.

Oracle Enterprise Manager returns to the Snapshots tab where you can see the new snapshot with its Parent Name listed in the table.

13.3 Managing Security Features with Oracle Enterprise Manager

This section describes how to manage security with Oracle Enterprise Manager.

- [Initializing Oracle ACFS Security](#) (page 13-7)
- [Preparing an Oracle ACFS For Security](#) (page 13-7)
- [Undoing Oracle ACFS Security](#) (page 13-7)
- [Enabling Security on Oracle ACFS](#) (page 13-8)
- [Disabling Security on Oracle ACFS](#) (page 13-8)
- [Viewing Security Status](#) (page 13-8)
- [Adding an Oracle ACFS Security Administrator](#) (page 13-9)
- [Changing an Oracle ACFS Security Administrator Password](#) (page 13-9)
- [Deleting an Oracle ACFS Security Administrator](#) (page 13-9)
- [Viewing Oracle ACFS Security Administrators](#) (page 13-9)
- [Creating a Rule](#) (page 13-10)
- [Cloning a Rule](#) (page 13-10)
- [Editing a Rule](#) (page 13-11)
- [Deleting a Rule](#) (page 13-11)
- [Searching For and Viewing Rules](#) (page 13-11)
- [Creating a Ruleset](#) (page 13-12)
- [Cloning a Ruleset](#) (page 13-12)
- [Editing a Ruleset](#) (page 13-13)
- [Deleting a Ruleset](#) (page 13-13)
- [Searching For and Viewing Rulesets](#) (page 13-13)
- [Creating a Realm](#) (page 13-14)
- [Cloning a Realm](#) (page 13-14)
- [Adding Objects to a Realm](#) (page 13-15)
- [Deleting Objects from a Realm](#) (page 13-15)
- [Deleting a Realm](#) (page 13-16)
- [Viewing Realms](#) (page 13-16)

For more information about Oracle ACFS security, refer to "[Oracle ACFS Security](#) (page 11-22)".

13.3.1 Initializing Oracle ACFS Security

Oracle ACFS security provides a unified security infrastructure that allows users to define security policies and groups for accessing file system objects. Oracle ACFS encryption allows users to store data on disk in an encrypted format. To initialize Oracle ACFS security, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.
For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".
2. Click a mount point in the list.
3. Click the **Security/Encryption** to display the Security and Encryption section.
4. In the **Security Configuration** section, you can click **Configuration, Rules, Rulesets, and Realms** under **Security** to display various options.
5. If Oracle ACFS security is not initialized in the cluster, the **Initialize Security** button displays on the page. Click the button to initialize Oracle ACFS security. This operation needs to be performed only once in the cluster. On this page, you can also prepare security for the file system.

13.3.2 Preparing an Oracle ACFS For Security

To prepare an Oracle ACFS for security, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.
For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".
2. Click the **Mount Point** of the file system for which you want to prepare security.
Enterprise Manager displays the ASM Cluster File System page for that mount point.
3. Click the **Security/Encryption** tab to display the Security Overview section. This page shows the security details for the current file system.
4. From the Actions menu, click **Prepare Security**.

13.3.3 Undoing Oracle ACFS Security

To undo preparation for Oracle ACFS security, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.
For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".
2. Click the **Mount Point** of the file system for which you want to unprepare security.
Enterprise Manager displays the ASM Cluster File System page for that mount point.

3. Click the **Security/Encryption** tab to display the Security Overview section. This page shows the security details for the current file system.
4. From the Actions menu, click **Unprepare Security**.

13.3.4 Enabling Security on Oracle ACFS

To enable security on Oracle ACFS, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".

2. Click the **Mount Point** of the file system for which you want to enable security.

Enterprise Manager displays the ASM Cluster File System page for that mount point.

3. Click the **Security/Encryption** tab to display the Security Overview section. This page shows the security details for the current file system.
4. From the Actions menu, click **Enable Security**.

13.3.5 Disabling Security on Oracle ACFS

To disable security on an Oracle ACFS, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".

2. Click the **Mount Point** of the file system for which you want to disable security.

Enterprise Manager displays the ASM Cluster File System page for that mount point.

3. Click the **Security/Encryption** tab to display the Security Overview section. This page shows the security details for the current file system.
4. From the Actions menu, click **Disable Security**.

13.3.6 Viewing Security Status

To view security status, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".

2. Click **Show Security and Encryption** to display the Security and Encryption section.

The section displays the mount points and the security and encryption settings for each.

13.3.7 Adding an Oracle ACFS Security Administrator

To add an Oracle ACFS security administrator, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click **Show Security and Encryption** to display the **Security and Encryption** section.
3. The list of Oracle ACFS security administrators in this cluster is displayed. Click the **Create** button. The **Create Administrator** page is displayed.
4. Enter the input details and click **OK**.

13.3.8 Changing an Oracle ACFS Security Administrator Password

To change an Oracle ACFS security administrator password, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click **Show Security and Encryption** to display the **Security and Encryption** section.
3. The list of Oracle ACFS security administrators in this cluster is displayed. Select a security administrator and click the **Change Password** button. The **Change Administrator Password** page is displayed.
4. Enter the input details and click **OK**.

13.3.9 Deleting an Oracle ACFS Security Administrator

To delete an Oracle ACFS security administrator, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click **Show Security and Encryption** to display the **Security and Encryption** section.
3. The list of Oracle ACFS security administrators in this cluster is displayed. Click the **Delete** button.

13.3.10 Viewing Oracle ACFS Security Administrators

To view Oracle ACFS security administrators, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click **Show Security and Encryption** to display the **Security and Encryption** section.
3. The list of Oracle ACFS security administrators in this cluster is displayed.

13.3.11 Creating a Rule

Rules are created in a file system and can be added to one or more rulesets, which can then be applied on an Oracle ACFS realm. Rules can also be cloned within the same file system or across file systems.

To create a rule, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Rules**.

The list of rules in the mount point is displayed.

5. Click **Create**. The Create Rule page displays along with the available rule options.
6. Enter the input details and click **OK**.

13.3.12 Cloning a Rule

Rules can be cloned within the same mount point or across mount points. To clone a rule, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Rules**.

The list of rules in the mount point is displayed.

5. Click **Clone**. The Clone Rule page is displayed with the available rule options.

6. Enter the input details and click **OK**.

13.3.13 Editing a Rule

You can modify values and options for an existing rules. Rule types cannot be changed. To edit a rule, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Rules**.

The list of rules in the mount point is displayed.

5. Select a rule and click **Edit**.
6. Edit the input details and click **OK**.

13.3.14 Deleting a Rule

You can delete a rule so it no longer is used to enforce security on realms. To delete a rule, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Rules**.

The list of rules in the mount point is displayed.

5. Select a rule and click the **Delete Rule** button.

13.3.15 Searching For and Viewing Rules

You can view or search the list of rules on a mount point, including the details of each rule. To search for and view rules, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.

3. Click the **Security/Encryption** tab.
4. In the **Security** section, click **Rules**.
The list of rules in the file system is displayed.

13.3.16 Creating a Ruleset

You can create a ruleset which is a collection of one or more rules that you can associate with a command rule for realm authorization. A ruleset evaluates to true or false based on the evaluation of each rule it contains and the evaluation type (All True or Any True). Rulesets cannot be applied on an Oracle ACFS security realm directly; they are always associated with a Command Rule, which are then applied on a realm to enforce security. To create a ruleset, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".
2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Rulesets**.

The list of rulesets in the mount point is displayed.
5. Click **Create**. The Create Ruleset page is displayed along with the available ruleset options.
6. Enter the input details and click **OK**.

13.3.17 Cloning a Ruleset

Rulesets can be cloned within the same mount point or across mount points. To clone a ruleset, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".
2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Rulesets**.

The list of rulesets in the mount point is displayed.
5. Click the ruleset you want to clone then click **Clone**. The Clone Ruleset page is displayed along with the available ruleset options.
6. Enter input details and click **OK**.

13.3.18 Editing a Ruleset

You can add or delete rules to or from a Ruleset. You can also edit the Ruleset options. To edit a ruleset, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Rulesets**.

The list of rulesets in the mount point is displayed.

5. Select a ruleset and click **Edit**. The Edit Ruleset page is displayed along with all the ruleset values.
6. Edit the ruleset details and click **OK**.

13.3.19 Deleting a Ruleset

Deleted rulesets are no longer used to enforce security on realms. To delete a ruleset, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page \(page 7-2\)](#)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Rulesets**.

The list of rulesets in the mount point is displayed.

5. Select a ruleset and click **Delete**.

13.3.20 Searching For and Viewing Rulesets

A ruleset is a collection of one or more rules that you can associate with a command rule for realm authorization. A ruleset evaluates to true or false based on the evaluation of each rule it contains and the evaluation type (All True or Any True). Enterprise Manager enables you to view and search the list of rulesets applied on an Oracle ACFS mount point, including the details of each ruleset. To search for and view rulesets, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the **Security** section, click **Rulesets**.
The list of rulesets in the file system is displayed.

13.3.21 Creating a Realm

Oracle ACFS Security realms provide the capability to classify file system objects into functional groups to provide fine-grained access control to data stored in these objects.

Realms are created at the file system level. To enforce security on the file objects in a realm, a ruleset is associated with a command rule which is then applied on the realm.

To create a realm, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Realms**.
5. Click **Create**. The Create Realm wizard displays along with the available realm options.
6. Enter the input details and click **OK**.

13.3.22 Cloning a Realm

Cloning a realm creates a new realm with all the objects of which the cloned realm consists. To clone a realm, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click a **Mount Point**. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Realms**. The list of realms in the mount point is displayed.
5. Select a realm and click **Clone Realm**. The **Clone Realm** page is displayed along with the available realm options. A realm can be cloned within the same mount

point with a different name or across mount points. You can choose the properties; such as users, groups, and encryption; to propagate to the cloned realm.

6. Enter the input details and click **OK**.

13.3.23 Adding Objects to a Realm

Once a realm is created, you can add objects to a realm. To add objects to a realm, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click the **Mount Point** that contains the realm to which you want to add objects. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Realms**. The list of realms in the mount point is displayed.
5. Select a realm and click the **Edit** button. The Edit Realm wizard displays.
6. Click **Next** to move to the Realm Objects page where you can click **Add** to add objects to the realm.
7. Click **Done**.

13.3.24 Deleting Objects from a Realm

You can delete various objects from a realm. If a file or directory is removed from a realm, it is not subjected to security or encryption that is enforced on the realm. If a command rule is deleted from a realm, it is not evaluated while enforcing security on the realm. The command rule still continues to exist for use in other realms. If an encryption object is removed, it disables encryption on the realm and results in decryption of all encrypted files in the realm.

To delete an object from a realm, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click the **Mount Point** that contains the realm from which you want to delete objects. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Realms**. The list of realms in the mount point is displayed.
5. Select a realm and click the **Edit** button. The Edit Realm wizard displays.
6. Click **Next** to move to the Realm Objects page where you can click **Remove** to delete objects from the realm.

7. Click **Done**.

13.3.25 Deleting a Realm

You can delete a realm from an Oracle ACFS file system. After a realm is deleted, access to all objects that were part of the realm is reverted to the state prior to their addition to the realm. These objects are not subjected to security and encryption that were enabled on the realm.

To delete a realm, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click the **Mount Point** that contains the realm you want to delete. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the Security section, click **Realms**. The list of realms in the mount point is displayed.
5. Select the realm you want to delete and click **Delete**.

13.3.26 Viewing Realms

To view realms, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Click the **Mount Point** that contains the realms you want to view. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. In the **Security** section, click **Realms**. The list of realms in the file system is displayed.

13.4 Managing Encryption Features with Oracle Enterprise Manager

This section describes how to manage encryption with Oracle Enterprise Manager.

- [Initializing Oracle ACFS Encryption](#) (page 13-17)
- [Enabling, Disabling, and Setting Parameters for Encryption On an Oracle ACFS](#) (page 13-17)
- [Viewing Encryption Status](#) (page 13-17)

For more information about Oracle ACFS encryption, refer to "[Oracle ACFS Encryption](#) (page 11-25)".

13.4.1 Initializing Oracle ACFS Encryption

To initialize Oracle ACFS encryption, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.
For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".
2. Click a mount point in the list.
3. Click the **Security/Encryption** tab to display the Security and Encryption section.
4. In the **Encryption Configuration** section, you can specify various options, such as **Unset Encryption Parameters**.
5. If Oracle ACFS encryption is not initialized in the cluster, the **Initialize Encryption** button displays. Click the button to initialize Oracle ACFS encryption. This operation needs to be performed only once in the cluster.

13.4.2 Enabling, Disabling, and Setting Parameters for Encryption On an Oracle ACFS

To enable, disable, or set parameters for Oracle ACFS encryption, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.
2. Click the **Mount Point** that contains the realms you want to view. The Oracle ACFS home page is displayed for this mount point.
3. Click the **Security/Encryption** tab.
4. Click **Encryption Configuration** link to display the Encryption configuration section.

13.4.3 Viewing Encryption Status

To view encryption status, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.
For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".
2. Click **Show Security and Encryption** to display the Security and Encryption section.

The section displays the mount points and the security and encryption settings for each.

13.5 Managing Tagging Features with Oracle Enterprise Manager

This section describes how to manage tagging with Oracle Enterprise Manager.

- [Adding a Tag in Oracle ACFS](#) (page 13-18)

- [Removing a Tag in Oracle ACFS](#) (page 13-18)
- [Searching for Tags in Oracle ACFS](#) (page 13-19)

For more information about Oracle ACFS tagging, refer to "[Oracle ACFS Tagging](#) (page 11-39)".

13.5.1 Adding a Tag in Oracle ACFS

You can add tags to directories and files that reside in a single Oracle ACFS file system or across multiple Oracle ACFS file systems. With the tag names, you can organize your files or perform operations, such as search and replication.

To add tags for specific files or directories from the ASM Cluster File System Page, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Select the file or directory you want to tag and select **Add Tag** from the Actions menu, then click **Go**.

Enterprise Manager displays the Add Tag page.

3. Add a Tag Name (or a comma-delimited list of names) in the Tag Name box. Optionally, you can turn on the option that enables you to add tags recursively to subdirectories and files.
4. In the Directories and Files section, enter either the directory paths directly as a comma-delimited list in the **Enter Directory** box, or search for directory paths with the ASM Cluster File System by selecting **Search Directory**. You can add a directory path by clicking **Add**.
5. Click **OK** to add tags.

13.5.2 Removing a Tag in Oracle ACFS

You can remove tags for specific files or directories in an Oracle ACFS file system. Optionally, you can recursively apply the remove operation to all subdirectories and files.

To remove tags, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. Select the file or directory from which you want to remove tags and select **Remove Tag** from the Actions menu, then click **Go**.

Enterprise Manager displays the Remove Tag page.

3. In the Tag Name box, enter the name of the tag you want to remove or specify a comma-delimited list of tag names to remove. Optionally, you can remove tags

recursively from all directories and files that exist in the specified directory paths by turning on the option.

4. In the Directories and Files section, select **Enter Directory** to enter the directory paths directly as a comma-delimited list, or select **Search Directory** to search for directory paths within an Oracle ACFS file system. To add directory paths repeatedly from different Oracle ACFS file systems, choose the file system and click **Add**.
5. Click **OK** to remove tags.

13.5.3 Searching for Tags in Oracle ACFS

You can search for tags from selected files and directories in an Oracle ACFS file system. Optionally, you can recursively apply this operation to all subdirectories and files. To search for tags for specific files or directories in an Oracle ACFS file system, follow these steps:

1. From the Oracle ASM home page, select the ASM Cluster File System tab to display the mount points in the cluster file system.

For information about the Oracle ASM home page and how to access that page, see "[Accessing the Oracle Automatic Storage Management Home Page](#) (page 7-2)".

2. From the Action menu, select **Search Tags**, then click **Go**.

Enterprise Manager displays the Search Tags page.

3. In the Search Criteria section, specify a comma-delimited list of tag names in the Tag Name box. Optionally, you can search tags recursively in all directories and files by turning on the option.
4. In the Directories and Files section, select **Enter Directory** to enter the directory path as a comma-delimited list, or select **Search Directory** to choose the Oracle ACFS file system and click **Add** to add directory paths repeatedly from different Oracle ACFS file systems.
5. Click **Search**.

The results of the search appear in the Search Results table.

6. Click **Return** to navigate back to the Oracle ASM home page displaying the ASM Cluster File System tab.

Managing Oracle ACFS and Oracle ADVM With ASMCA

Oracle ASM Configuration Assistant (ASMCA) provides utilities for managing Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM).

Oracle ASM Configuration Assistant (ASMCA) supports installing and configuring Oracle ASM instances, Oracle ASM disk groups, Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes, and Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems. In addition, you can use the ASMCA command-line interface.

This chapter discusses the following topics:

- [ASMCA GUI Tool for Managing Oracle ACFS and Oracle ADVM](#) (page 14-1)
- [ASMCA Command-Line Interface for Managing Oracle ACFS and Oracle ADVM](#) (page 14-6)

For information about managing Oracle ASM instances and disk groups with ASMCA, refer to [Managing Oracle ASM With ASMCA](#) (page 9-1).

Note:

Some ASMCA functionality and menu options are only available in specific environments, such as an Oracle RAC configuration. For details about a specific ASMCA page, access the online help provided with the **Help** button.

14.1 ASMCA GUI Tool for Managing Oracle ACFS and Oracle ADVM

This section contains the following topics:

- [Managing Oracle ADVM Volumes with ASMCA](#) (page 14-2)
- [Managing Oracle ACFS File Systems with ASMCA](#) (page 14-3)
- [Managing Security and Encryption for Oracle ACFS with ASMCA](#) (page 14-4)
- [Creating an Oracle ACFS File System for Database Use](#) (page 14-5)

14.1.1 Managing Oracle ADVM Volumes with ASMCA

Note:

When creating an Oracle ADVM volume for an Oracle ACFS file system that is intended to store database files, use the ASMCMD `volcreate` command or the SQL `ALTER DISKGROUP ADD VOLUME` SQL statement rather than the ASMCA tool to ensure that the column striping is set to 1.

Oracle ASM Configuration Assistant enables you to create or configure an Oracle ADVM volume.

Right click a selected volume in the Oracle ASM Configuration Assistant screen to display the configuration options menu.

The menu includes options to:

- Enable disabled volumes and disable enabled volumes
- View status details
- Resize volumes that have not been mounted on an Oracle ACFS file system
- Delete volumes

There are options for creating a volume, enabling all volumes, and disabling all volumes.

Click **Create** to display the dialog for creating a volume. You must provide a unique volume name for the existing Oracle ASM disk group that you select from the list of disk groups. You must also 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 to contain an Oracle ADVM volume. To use Oracle ACFS encryption, replication, security, or tagging, the disk group on which the volume is created for the file system must have compatibility attributes for ASM and ADVM set to 11 . 2 . 0 . 2 or higher.

Before creating an Oracle ADVM volume on AIX, ensure that the necessary user authorizations have been created.

See Also:

- [Managing Oracle ADVM Volumes in a Disk Group](#) (page 4-21) for information about the ALTER DISKGROUP ADD VOLUME statement
 - [volcreate](#) (page 15-2) for information about the advanced options when creating a volume
 - [Disk Group Compatibility Attributes](#) (page 4-51) for information about disk group compatibility
 - [Oracle ACFS Command-Line Tools for the AIX Environment](#) (page 16-41) for information about creating an Oracle ADVM volume on the AIX operating system
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about creating a volume with the ASMCMD volcreate command
 - [Introducing Oracle ACFS and Oracle ADVM](#) (page 11-1) for an overview of Oracle ACFS
-

14.1.2 Managing Oracle ACFS File Systems with ASMCA

Oracle ASM Configuration Assistant enables you to create or configure an Oracle ACFS file system.

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 includes options to:

- View status details
- Mount and dismount the file system
- Resize the file system
- Delete the file system
- Display or execute the security commands
- Create or delete snapshots

Some commands require root privileges, such as mounting a file system. ASMCA generates the command for you to run manually as root.

When creating Oracle ACFS file systems on Windows, run ASMCA as a Windows domain user who is also an administrator of the computer.

There are buttons for Create, Mount All, Dismount All, and Security and Encryption commands.

Oracle ASM Configuration Assistant enables you to create an Oracle ACFS file system.

Select **Create**, then select the type of Oracle ACFS file system you want to create and an existing Oracle ADVM volume from the lists. Specify the mount point field; the mount point must be an existing directory. The file system must be mounted to make it available. You can choose to enable the **Auto Mount** option. You can also provide a description of the file system in the **Description** field.

The mount command must be manually run as root or the Windows Administrator at an operating system prompt.

You can choose to create an Oracle ACFS file system for an Oracle Database.

See Also:

- [acfsmountvol](#) (page 16-54) for information about mounting an Oracle ACFS file system on Windows
 - [mount](#) (page 16-30) for information about mounting an Oracle ACFS file system on Linux
 - [Creating an Oracle ACFS File System for Database Use](#) (page 14-5) for more information about creating a file system for a database home
 - [Introducing Oracle ACFS and Oracle ADVM](#) (page 11-1) for additional information about Oracle ACFS
 - [Basic Steps to Manage Oracle ACFS Systems](#) (page 16-2) for a summary of the basic steps for creating an Oracle ACFS file system
-
-

14.1.3 Managing Security and Encryption for Oracle ACFS with ASMCA

Oracle ASM Configuration Assistant enables you to configure security and encryption for an Oracle ACFS file system.

You must initialize the security system as the first step in configuring security for an Oracle ACFS file system. You must also initialize the encryption system as the first step in encrypting an Oracle ACFS file system. You do not have to use both Oracle ACFS security and encryption on the same file system. If you decide to use both security and encryption, then encryption must be initialized and set before enabling encryption on a security realm.

In the dialog that displays, enter the information to specify the security administrator and the operating system group of the security administrator. You can choose to create a password protected wallet. After you complete the entry fields in the dialog, click **Show Command** to display the commands you must run as a root or Administrator user at an operating system prompt. For example:

```
# /sbin/acfsutil sec init -u grid -g asmadmin
# /sbin/acfsutil encr init
```

After security has been initialized, you can use the menu options to manage security and encryption for an Oracle ACFS file system.

The menu includes options to:

- Set encryption
- Enable and disable encryption
- Prepare and enable security
- Enable and disable security

See Also:

- - [acfsutil sec init](#) (page 16-93) and [acfsutil encr init](#) (page 16-119) for information about the `acfsutil sec init` and `acfsutil encr init` commands
 - [Oracle ACFS Security](#) (page 11-22) for information about Oracle ACFS security
 - [Securing Oracle ACFS File Systems](#) (page 16-9) for a summary of the basic steps for setting up security on an Oracle ACFS file system
 - [Oracle ACFS Encryption](#) (page 11-25) for information about Oracle ACFS encryption
 - [Encrypting Oracle ACFS File Systems](#) (page 16-14) for a summary of the basic steps for setting up encryption on an Oracle ACFS file system
 - [Oracle ACFS Command-Line Utilities](#) (page 16-143) for information about the `acfsutil encr` and `acfsutil sec` commands
-
-

14.1.4 Creating an Oracle ACFS File System for Database Use

To create an Oracle ACFS file system for database use, you can select the **Create ACFS for Database Use** from the disk group configuration options menu to display the dialog entry box.

Note:

When creating an Oracle ADVM volume for an Oracle ACFS file system that is intended to store database files, use the `ASMCMD volcreate` command or the `SQL ALTER DISKGROUP ADD VOLUME SQL` statement rather than the ASMCA tool to ensure that the column striping is set to 1.

In this dialog, you must enter:

- **Volume Name**
This is the name of the Oracle ADVM volume you want to create.
- **Mount Point**
This is 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.
- **Size in gigabytes (GB)**
The default is 7 GB and the minimum recommended size.
- **Owner Name**
This is the operating system name of the user that installs the database and owns the software in the 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.

Select **Automatically run configuration commands** to run ASMCA configuration commands automatically. To use this option, you must provide the root credentials on the ASMCA Settings page.

The mount command may also be run manually as root or the Windows Administrator at an operating system prompt.

See Also:

- [volcreate](#) (page 15-2) for information about the `volcreate` command
 - [Managing Oracle ADVM Volumes in a Disk Group](#) (page 4-21) for information about the `ALTER DISKGROUP ADD VOLUME` statement
 - [Overview of Oracle ASM Dynamic Volume Manager](#) (page 11-49) for information about Oracle ADVM volumes
 - [About Oracle ACFS and Oracle Database Homes](#) (page 11-8) for information about mount points and database homes
 - [mount](#) (page 16-30) for information about mounting an Oracle ACFS file system on Linux
 - [acfsmountvol](#) (page 16-54) for information about mounting an Oracle ACFS file system on Windows
 - [About the Oracle ACFS Mount Registry](#) (page 11-11) for information on the mount registry
-
-

14.2 ASMCA Command-Line Interface for Managing Oracle ACFS and Oracle ADVM

The ASMCA command-line interface provides non-GUI support for configuring Oracle ASM disk groups, volumes, and Oracle ACFS.

- [ASMCA Commands for Oracle ACFS and Oracle ADVM](#) (page 14-6)

For general information about running ASMCA command-line interface, refer to "[Running ASMCA Command-Line](#) (page 9-4)".

14.2.1 ASMCA Commands for Oracle ACFS and Oracle ADVM

This section describes the commands that can be run with ASMCA command-line to manage Oracle ACFS objects.

- [Create an Oracle ACFS Snapshot](#) (page 14-7)
- [Delete an Oracle ACFS Snapshot](#) (page 14-7)
- [Create a Volume](#) (page 14-8)

- [Create an Oracle ACFS File System](#) (page 14-8)

Run ASMCA commands as a Windows domain user.

14.2.1.1 Create an Oracle ACFS Snapshot

`-createACFSSnapshot` creates an Oracle Automatic Storage Management Cluster File System (Oracle ACFS) snapshot.

Syntax

```
asmca -silent
  -createACFSSnapshot
    (-acfsMountPoint mount_path )
    (-snapshotName snapshot_name )
    [-snapshotMode { r|w } ]
    [-parentSnapshotName parent_snapshot_name ]
```

[Table 14-1](#) (page 14-7) contains the options available with the `-createACFSSnapshot` command.

Table 14-1 Options for the `-createACFSSnapshot` command

Option	Description
<code>-acfsMountPoint mount_path</code>	Specifies an existing path to be used as the mount point. This is required in an Oracle RAC environment.
<code>-snapshotName snapshot_name</code>	Specifies the name of the Oracle ACFS snapshot.
<code>-snapshotMode r w</code>	Specifies the mode of the Oracle ACFS snapshot, either read-only (r) or read-write (w).
<code>-parentSnapshotName parent_snapshot_name</code>	Specifies the name of the Oracle ACFS parent snapshot.

14.2.1.2 Delete an Oracle ACFS Snapshot

`-deleteACFSSnapshot` deletes an existing Oracle Automatic Storage Management Cluster File System (Oracle ACFS) snapshot.

Syntax

```
asmca -silent
  -deleteACFSSnapshot
    (-acfsMountPoint mount_path )
    (-snapshotName snapshot_name )
```

[Table 14-2](#) (page 14-7) contains the options available with the `-deleteACFSSnapshot` command.

Table 14-2 Options for the `-deleteACFSSnapshot` command

Option	Description
<code>-acfsMountPoint mount_path</code>	Specifies an existing path to be used as the mount point. This is required in an Oracle RAC environment.

Table 14-2 (Cont.) Options for the `-deleteACFSSnapshot` command

Option	Description
<code>-snapshotName <i>anapshot_name</i></code>	Specifies the name of the Oracle ACFS snapshot.

14.2.1.3 Create a Volume

`-createVolume` creates Oracle ADVM volumes.

Syntax

```
asmca -silent
      -createVolume
        { -volumeName volume_name
          -volumeDiskGroup diskgroup
          -volumeSizeGB size_GB
          [ -volumeRedundancy { INHERIT | HIGH | NORMAL | EXTERNAL } ] ... }
      [-sysAsmPassword sysasm_password ]
```

Table 14-3 (page 14-8) contains the options available with the `-createVolume` command.

Table 14-3 Options for the `-createVolume` command

Option	Description
<code>-volumeName <i>volume_name</i></code>	Specifies the name of the volume to create.
<code>-volumeDiskGroup <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 9-1](#) (page 9-4).

Examples

To create an Oracle ADVM volume:

Example 14-1 Using `asmca -silent -createVolume`

```
$ asmca -silent
      -createVolume
        -volumeName volume1
        -volumeDiskGroup mynewdg
        -volumeSizeGB 1
```

Volume *volume1* created successfully.

14.2.1.4 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](#) (page 16-30)" for Linux environments and "[acfsmountvol](#) (page 16-54)" for Windows environments.

Syntax

```
asmca -silent
      -createACFS
        { -acfsVolumeDevice volume_device }
        { -acfsMountPoint mount_path }
        [-acfsUserName acfs_username ]
        [-acfsUserGroup acfs_usergroup ]
        [-sysAsmPassword sysasm_password ]
```

[Table 14-4](#) (page 14-9) contains the options available with the `-createACFS` command.

Table 14-4 Options for the `-createACFS` command

Option	Description
<code>-acfsVolumeDevice <i>volume_device</i></code>	Specifies the name of the Oracle ADVM volume device.
<code>-acfsMountPoint <i>mount_path</i></code>	Specifies an existing path to be used as the mount point. This is required in an Oracle RAC environment.
<code>-acfsUserName <i>acfs_username</i></code>	Specifies the Oracle ACFS user name.
<code>-acfsUserGroup <i>acfs_usergroup</i></code>	Specifies the Oracle ACFS group name.

For additional options that are common to multiple commands, see [Table 9-1](#) (page 9-4).

Examples

To create an Oracle ACFS file system:

Example 14-2 Using `asmca -silent -createACFS`

```
$ asmca -silent
      -createACFS
        -acfsVolumeDevice /dev/asm/volume1-457
        -acfsMountPoint /acfsmounts/acfs1
```

Managing Oracle ADVM with ASMCMD

Oracle ASM Command-Line Utility (ASMCMD) provides commands for managing Oracle ADVM volumes.

This chapter describes the Oracle Automatic Storage Management (Oracle ASM) Command-Line Utility (ASMCMD) volume management commands.

- [ASMCMD Volume Management Commands](#) (page 15-1)

See Also:

- [Overview of Oracle ASM Dynamic Volume Manager](#) (page 11-49) for information about Oracle ASM Dynamic Volume Manager (Oracle ADVM)
 - [Managing Oracle ADVM Volumes in a Disk Group](#) (page 4-21) for information about using SQL statements to manage Oracle ADVM volumes
 - [About ASMCMD](#) (page 10-1) for information about using ASMCMD
-

15.1 ASMCMD Volume Management Commands

This topic provides a summary of the Oracle ADVM volume management commands.

[Table 15-1](#) (page 15-1) lists the Oracle ADVM volume management commands with brief descriptions. To successfully run 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 15-1 Summary of ASMCMD volume management commands

Command	Description
volcreate (page 15-2)	Creates an Oracle ADVM volume in the specified disk group.
voldelete (page 15-4)	Deletes an Oracle ADVM volume.
voldisable (page 15-5)	Disables Oracle ADVM volumes in mounted disk groups.
volenable (page 15-6)	Enables Oracle ADVM volumes in mounted disk groups.
volinfo (page 15-6)	Displays information about Oracle ADVM volumes.
volresize (page 15-7)	Resizes an Oracle ADVM volume.
volset (page 15-8)	Sets attributes of an Oracle ADVM volume in mounted disk groups.

Table 15-1 (Cont.) Summary of ASMCMD volume management commands

Command	Description
volstat (page 15-9)	Reports volume I/O statistics.

15.1.1 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 15-2](#) (page 15-2) describes the options for the `volcreate` command.

Table 15-2 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 value must be a positive integer. The unit designation must be appended to the number specified. A space is not allowed between the number and the unit designation. For example: 200M or 20G
<code>--column <i>number</i></code>	Number of columns in a stripe set. Values range from 1 to 8. The default value is 8.
<code>--width <i>stripe</i></code>	Stripe width of a volume in units of K or M. The value can range from 4 KB to 1 MB, at power-of-two intervals. The default is 1M.
<code>--redundancy {<i>high</i> <i>mirror</i> <i>unprotected</i>}</code>	Redundancy of the Oracle ADVM volume which can be only specified for normal redundancy disk groups. If redundancy is not specified, the setting defaults to the redundancy level of the disk group and this is the recommended setting. The range of values is as follows: <code>unprotected</code> for non-mirrored redundancy, <code>mirror</code> for double-mirrored redundancy, or <code>high</code> for triple-mirrored redundancy.
<code>--primary {<i>hot</i> <i>cold</i>}</code>	Intelligent Data Placement specification for primary extents, either <code>hot</code> or <code>cold</code> region.
<code>--secondary {<i>hot</i> <i>cold</i>}</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. Only alphanumeric characters and underscores are allowed. Hyphens are not allowed. The first character must be alphabetic.

WARNING:

Specifying `--redundancy unprotected` means that Oracle ASM mirroring is not available for data recovery with the Oracle ADVM volume. The redundancy setting (`normal`) of the disk group does not provide mirroring for an unprotected Oracle ADVM volume. The `unprotected` configuration is not recommended for production environments as intermittent storage access failures can result in the loss of data. Backups are strongly recommended.

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`, such as `volume1-123`. On Windows the volume device name is in the format `asm-volume_name-nnn`, such as `asm-volume1-123`.

On Linux platforms, the volume name must be less than or equal to eleven alphanumeric characters, starting with an alphabetic character. On AIX platforms, the volume name must be less than or equal to twenty three alphanumeric characters, starting with an alphabetic character. On Windows and Solaris platforms, the volume name must be less than or equal to thirty alphanumeric characters, starting with an alphabetic character.

You can determine the volume device name with the `volinfo` command.

If the `--column` option is set to 1, then striping is disabled and the stripe width equals the volume extent size. The volume extent size is 8 MB if the disk group allocation unit (AU) is less than or equal to 8 MB. If the AU size is greater than 8 MB, then the Oracle ADVM volume extent size is equivalent to the disk group AU size. Setting the `--column` option to 8 (the default) is recommended to achieve optimal performance with database data files and other files.

A successful volume creation automatically enables the volume device.

The volume device file functions as any other disk or logical volume to mount file systems or for applications to use directly.

When creating an accelerator volume, create the volume on a disk group with storage that is significantly faster than the primary volume's storage.

Before creating an Oracle ADVM volume on AIX, ensure that the necessary user authorizations have been created.

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 15-1 Using the ASMCMDB volcreate command

```
ASMCMDB [+] > volcreate -G data -s 10G --width 1M --column 8 volume1
```

```
ASMCMDB [+] > volinfo -G data volume1
Diskgroup Name: DATA
```

```
Volume Name: VOLUME1
Volume Device: /dev/asm/volume1-123
State: ENABLED
```

Size (MB): 10240
Resize Unit (MB): 64
Redundancy: MIRROR
Stripe Columns: 8
Stripe Width (K): 1024
Usage:
Mountpath:

See Also:

- [Creating an Oracle ACFS File System](#) (page 16-5) for information on mounting the volume device file
 - [volinfo](#) (page 15-6) for information about the `volinfo` command.
 - [Limits of Oracle ADVM](#) (page 17-4) for information about Oracle ADVM limits
 - [mkfs](#) (page 16-28) for information about the accelerator volume
 - [Mirroring, Redundancy, and Failure Group Options](#) (page 4-38) for information about redundancy settings
 - [Intelligent Data Placement](#) (page 4-26) for information about hot and cold disk regions
 - [Oracle ACFS Command-Line Tools for the AIX Environment](#) (page 16-41) for information about user authorizations on AIX
-
-

15.1.2 voldelete

Purpose

Deletes an Oracle ADVM volume.

Syntax and Description

```
voldelete -G diskgroup volume
```

[Table 15-3](#) (page 15-4) describes the options for the `voldelete` command.

Table 15-3 Options for the `voldelete` command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the disk group containing the volume.
<code><i>volume</i></code>	Name of the volume.

To successfully run 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 15-2 Using the ASMCMD `voldelete` command

```
ASMCMD [+] > voldelete -G data volume1
```

15.1.3 `voldisable`

Purpose

Disables Oracle ADVM volumes in mounted disk groups and removes the volume device on the local node.

Syntax and Description

```
voldisable { --all | { -G diskgroup { -a | volume } } }
```

[Table 15-4](#) (page 15-5) describes the options for the `voldisable` command.

Table 15-4 Options for the `voldisable` command

Option	Description
<code>--all</code>	Specifies all volumes within all disk groups.
<code>-G <i>diskgroup</i></code>	Name of the disk group containing the volume.
<code>-a</code>	Specifies all volumes within the specified disk group.
<code><i>volume</i></code>	Name of the volume.

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 because 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](#)" (page 16-24).

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 15-3 Using the ASMCMD `voldisable` command

```
ASMCMD [+] > voldisable -G data volume1
```

15.1.4 volenable

Purpose

Enables Oracle ADVM volumes in mounted disk groups.

Syntax and Description

```
volenable { --all |{ -G diskgroup { -a | volume } } }
```

Table 15-5 (page 15-6) describes the options for the volenable command.

Table 15-5 Options for the volenable command

Option	Description
--all	Specifies all volumes within all disk groups.
-G <i>diskgroup</i>	Name of the disk group containing the volume.
-a	Specifies all volumes within the specified disk group.
<i>volume</i>	Name of the volume.

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 15-4 Using the ASMCMD volenable command

```
ASMCMD [+] > volenable -G data volume1
```

15.1.5 volinfo

Purpose

Displays information about Oracle ADVM volumes.

Syntax and Description

```
volinfo { --all |{ -G diskgroup { -a | volume } } }  
volinfo {--show_diskgroup |--show_volume} volumedev
```

Table 15-6 (page 15-6) describes the options for the volinfo command.

Table 15-6 Options for the volinfo command

Option	Description
--all	Specifies all volumes within all disk groups.
-G <i>diskgroup</i>	Name of the disk group containing the volume.
-a	Specifies all volumes within the specified disk group.

Table 15-6 (Cont.) Options for the volinfo command

Option	Description
<i>volume</i>	Name of the volume.
<code>--show_diskgroup</code>	Returns only the disk group name. A volume device name is required.
<code>--show_volume</code>	Returns only the volume name. A volume device name is required.
<i>volumedevic</i>	Name of the volume device.

Examples

The first example displays information about the `volume1` volume in the `data` disk group. This example was produced in a Linux environment. The second example displays information about the `volume1` volume in the `data` disk group and was produced in a Windows environment.

The `Mountpath` field contains the path where the volume is currently mounted or where the volume was last mounted.

Example 15-5 Using the ASMCMD volinfo command

```

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): 64
  Redundancy: MIRROR
  Stripe Columns: 8
  Stripe Width (K): 1024
  Usage: ACFS
  Mountpath: /acfsmounts/acfs1

ASMCMD [+] > volinfo -G data -a
Diskgroup Name: DATA
  Volume Name: VOLUME1
  Volume Device: \\.\asm-volume1-311
  State: ENABLED
  Size (MB): 10240
  Resize Unit (MB): 64
  Redundancy: MIRROR
  Stripe Columns: 8
  Stripe Width (K): 1024
  Usage: ACFS
  Mountpath: C:\oracle\acfsmounts\acfs1
    
```

15.1.6 volresize

Purpose

Resizes an Oracle ADVM volume.

Syntax and Description

```
volresize -G diskgroup -s size [ -f ] volume
```

[Table 15-7](#) (page 15-8) describes the options for the `volresize` command.

Table 15-7 Options for the `volresize` command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the disk group containing the volume.
<code>-f</code>	Forces the shrinking of a volume that is not an Oracle ACFS volume and suppresses any warning message.
<i>volume</i>	Name of the volume.
<code>-s <i>size</i></code>	New size of the volume in units of K, M, G, or T.

If the volume is mounted on a non-Oracle ACFS file system, then dismount the file system first before resizing. If 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, then 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](#) (page 16-164)".

Examples

The following is an example of the `volresize` command that resizes `volume1` in the `data` disk group to 20 gigabytes.

Example 15-6 Using the ASMCMD `volresize` command

```
ASMCMD [+] > volresize -G data -s 20G volume1
```

15.1.7 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 15-8](#) (page 15-8) describes the options for the `volset` command.

Table 15-8 Options for the `volset` command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the disk group containing the volume.

Table 15-8 (Cont.) Options for the volset command

Option	Description
<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 an Oracle ACFS 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 set when the file system is mounted and should not be changed.
<code>--primary {hot cold}</code>	Intelligent Data Placement specification for primary extents, either hot or cold region.
<code>--secondary {hot cold}</code>	Intelligent Data Placement specification for secondary extents, either hot or cold region.
<i>volume</i>	Name of the volume.

When running the `mkfs` command to create a file system, the `usage` field is set to ACFS and `mountpath` field is reset to an empty string if it has been set. The `usage` field should remain at ACFS.

When running the `mount` command to mount a file system, the `mountpath` field is set to the mount path value to identify the mount point for the file system. After the value is set by the `mount` command, the `mountpath` field should not be updated.

Examples

The following is an example of a `volset` command that sets the usage string for a volume that is not associated with a file system.

Example 15-7 Using the ASMCMD volset command

```
ASMCMD [+] > volset -G data --usagestring 'no file system created' volume1
```

15.1.8 volstat

Purpose

Reports I/O statistics for Oracle ADVM volumes.

Syntax and Description

```
volstat [-G diskgroup] [volume]
```

[Table 15-9](#) (page 15-9) describes the options for the `volstat` command.

Table 15-9 Options for the volstat command

Option	Description
<code>-G <i>diskgroup</i></code>	Name of the mounted disk group containing the volume.
<i>volume</i>	Name of the volume.

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 15-8 *Using the ASMCMD volstat command*

```
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
```

Managing Oracle ACFS with Command-Line Tools

Numerous command-line tools are provided for managing Oracle ASM Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM).

Command-line tools for managing and implementing Oracle Automatic Storage Management Cluster File System (Oracle ACFS), include the following:

- [Basic Steps to Manage Oracle ACFS Systems](#) (page 16-2)
- [Oracle ACFS Command-Line Tools for Linux Environments](#) (page 16-25)
- [Oracle ACFS Command-Line Tools for the Solaris Environment](#) (page 16-33)
- [Oracle ACFS Command-Line Tools for Windows Environments](#) (page 16-49)
- [Oracle ACFS Command-Line Tools for Tagging](#) (page 16-60)
- [Oracle ACFS Command-Line Tools for Replication](#) (page 16-63)
- [Oracle ACFS Command-Line Tools for Auditing](#) (page 16-81)
- [Oracle ACFS Command-Line Tools for Security](#) (page 16-86)
- [Oracle ACFS Command-Line Tools for Encryption](#) (page 16-117)
- [Oracle ACFS Command-Line Tools for Snapshots](#) (page 16-125)
- [Oracle ACFS Command-Line Tools for Compression](#) (page 16-140)
- [Oracle ACFS Command-Line Utilities](#) (page 16-143)

Note:

Oracle does not recommend using identifiers for Oracle Database object names that must be quoted. While the use of quoted identifiers may be valid as names in some command-line tools or in SQL statements, the names may not be valid when using other tools that manage the object.

See Also:

- [About Using Oracle ACFS Command-Line Tools](#) (page 16-2) for information about running Oracle ACFS `acfsutil` commands
 - *Oracle Database SQL Language Reference* for more information about naming Oracle Database objects
-
-

16.1 Basic Steps to Manage Oracle ACFS Systems

This topic provides an overview of the basic steps when managing Oracle ACFS file systems using command-line utilities.

The examples in this section show operating system commands that are run in a Linux environment system. ASMCMD commands manage the Oracle ADVM volumes, but you can also use SQL*PLUS and Oracle ASM Configuration Assistant (ASMCA) to manage volumes.

This section contains these topics:

- [About Using Oracle ACFS Command-Line Tools](#) (page 16-2)
- [Creating an Oracle ACFS File System](#) (page 16-5)
- [Accessing an Oracle ACFS File System on a Different Node in the Cluster](#) (page 16-7)
- [Managing Oracle ACFS Snapshots](#) (page 16-8)
- [Securing Oracle ACFS File Systems](#) (page 16-9)
- [Encrypting Oracle ACFS File Systems](#) (page 16-14)
- [Tagging Oracle ACFS File Systems](#) (page 16-16)
- [Replicating Oracle ACFS File Systems](#) (page 16-17)
- [Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#) (page 16-23)
- [Removing an Oracle ACFS File System and a Volume](#) (page 16-24)

16.1.1 About Using Oracle ACFS Command-Line Tools

This topic provides an overview of the use of Oracle ACFS `acfsutil` commands.

The discussions include:

- Privileges to Run Oracle ACFS `acfsutil` Commands
- Displaying Help for Oracle ACFS `acfsutil` Commands
- Running Oracle ACFS `acfsutil` Commands on Windows
- Displaying Oracle ACFS Version Information
- Managing Trace File Space for `acfsutil` Commands

Privileges to Run Oracle ACFS `acfsutil` Commands

To run many Oracle ACFS `acfsutil` commands, you must be a system administrator or an Oracle ASM administrator user that has been enabled to run the commands. These privileges are described as follows:

On Non-Windows systems:

- For system administrator privileges, you must be the `root` user.
- For Oracle ASM administrator user privileges, you must belong to the `OSASM` group and the `oinstall` group (for the `OINSTALL` privilege).

On Windows systems:

- For system administrator privileges, you must belong to the Administrators group.
- For Oracle ASM administrator user privileges, you must belong to the `ORA_ASMADMIN` group and the `ORA_CRS_USERS` group.

Displaying Help for Oracle ACFS `acfsutil` Commands

You can display help and usage text for Oracle ACFS `acfsutil` commands with the `h` option. When you include a command or a subcommand with the command, the help and usage display is specific to the command and subcommand entered.

The following example illustrates several different ways to display help and usage text, from the most general to more specific. This example shows the `-h` format to display help on a non-Windows platform. On Windows, use `/h` in place of `-h`.

Example 16-1 *Displaying help for Oracle ACFS `acfsutil` commands*

```
$ /sbin/acfsutil -h

$ /sbin/acfsutil -h compress
$ /sbin/acfsutil compress -h

$ /sbin/acfsutil -h repl info
$ /sbin/acfsutil repl info -h

$ /sbin/acfsutil -h sec admin info
$ /sbin/acfsutil sec admin info -h
```

Running Oracle ACFS `acfsutil` Commands on Windows

When the options are specified with Oracle ACFS `acfsutil` commands on a Windows platform, use `/` in place of `-` with the option. For example, you can display help for `acfsutil` commands on a Linux platform with `acfsutil -h`. On a Windows platform, use `acfsutil /h`.

A mount point on a Windows operating system can be a just the drive letter (`M:`) or a directory including the drive letter (`M:\my_mount_point`).

When an `acfsutil` command on Windows targets the root of the file system that is mounted on a drive letter, include the backslash and a period (`\.`) with the drive letter (`P:\.`) to avoid the possibility of triggering a Windows path substitution to the last accessed path on the specified drive. For example:

```
C:\oracle> acfsutil info fs P:\.
```

Displaying Oracle ACFS Version Information

You can run `acfsutil version` to display the Oracle ACFS version. For example:

```
$ /sbin/acfsutil version
acfsutil version: 12.2.0.0.3
```

Managing Trace Files for `acfsutil` Commands

The Automatic Diagnostic Repository (ADR) generates a separate internal file for each `acfsutil` command invocation to trace the operation of the command. The space consumed by these trace files can increase significantly, and some features, such as snapshot-based replication, may generate a significant number of trace files.

To limit the number of trace files and the space consumed by them, you can set policy attributes with the Automatic Diagnostic Repository Command Interpreter (ADRCI) utility to purge trace files after a specified retention period. ADRCI considers trace files to be short-lived files and the retention period is controlled by the setting of the `SHORTP_POLICY` attribute. You can view the current retention period for these trace files with the ADRCI `show control` command.

By default, the short-lived files are retained for 720 hours (30 days). The value in hours specifies the number of hours after creation when a given file is eligible for purging. To limit the number of these files and the space consumed by them, you can update the number of hours set for the `SHORTP_POLICY` retention period, such as 240 hours (10 days).

The following steps summarize how to update the retention period for short-lived trace files.

- Start the Automatic Diagnostic Repository Command Interpreter (ADRCI) utility.

```
$ adcri
```

- Display the ADR home directory paths (ADR homes):

```
ADRCI> show homes
```

- If more than one home is shown, then set the appropriate home for the trace files you want to administer:

```
ADRCI> set homepath my_specified_homepath
```

- Display the current configuration values.

```
ADRCI> show control
```

- Update a specific ADRCI configuration value. For example, set `SHORTP_POLICY` to 240 hours (10 days).

In the displayed `show control` output, check the value of the `SHORTP_POLICY` attribute, which is the retention period in hours for short-lived files. If necessary, set a new retention period for short-lived trace files with the following:

```
ADRCI> set control (SHORTP_POLICY=240)
```

If you want to start an immediate purge of the trace files in the current ADR home path, you can use the following command:

```
ADRCI> purge -type TRACE -age number_of_minutes
```

The value *number_of_minutes* controls which files are purged based on the age of the files. Files older than the specified number of minutes are targeted for the purge operation.

See Also:

- [About Privileges for Oracle ASM](#) (page 3-28) for information about user privileges for Oracle ASM
 - *Oracle Database Utilities* for information about the Automatic Diagnostic Repository Command Interpreter (ADRCI) utility
-
-

16.1.2 Creating an Oracle ACFS File System

You can create an Oracle ACFS file system using the steps in this topic.

To create and verify a file system, perform the following steps:

1. Create an Oracle ADVM 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 to contain an Oracle ADVM volume. To use Oracle ACFS encryption, replication, security, or tagging, the disk group on which the volume is created for the file system must have compatibility attributes for ASM and ADVM set to 11.2.0.2 or higher.

Start ASMCMD connected to the Oracle ASM instance. You must be a user in the OSASM operating system group.

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
```

When creating an Oracle ADVM volume, a volume device name is created that includes a unique Oracle ADVM persistent disk group number. The volume device file functions in the same manner as any other disk or logical volume to mount file systems or for applications to use directly.

The format of the volume name is platform-specific.

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

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.
```

The `root` privilege is not required to run `mkfs`. The ownership of the volume device file dictates who can run this command.

4. Optionally register the file system with the `acfsutil registry` command.

For example:

```
$ /sbin/acfsutil registry -a /dev/asm/volume1-123 /acfsmounts/acfs1

acfsutil registry: mount point /acfsmounts/acfs1 successfully added
to Oracle Registry
```

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 requirement 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.

Note:

Oracle ACFS registration (`acfsutil registry`) is not supported in an Oracle Restart (standalone) 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 /acfsmounts/acfs1
```

The `root` privilege is required to run the `mount` command and the Windows Administrator privilege is required to run the `acfsmountvol` command.

After the file system has been mounted, ensure that the permissions are set to allow access to the file system for the appropriate users. For example:

```
# chown -R oracle:dba /acfsmounts/acfs1
```

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" > /acfsmounts/acfs1/myfile
```

7. List the contents of the test file that was created in the file system.

For example:

```
$ cat /acfsmounts/acfs1/myfile
Oracle ACFS File System
```

See Also:

- [Disk Group Compatibility Attributes](#) (page 4-51) for information about disk group compatibility settings.
 - [About Privileges for Oracle ASM](#) (page 3-28) for information about operating system groups and privileges
 - [mkfs](#) (page 16-28) (Linux environments) or [acfsformat](#) (page 16-53) (Windows) for information about commands to create an Oracle ACFS file system.
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about the `volcreate` command and the `volinfo` command
 - [acfsutil registry](#) (page 16-160) for information about the `acfsutil registry` command to register an Oracle ACFS file system
 - [About the Oracle ACFS Mount Registry](#) (page 11-11) for about information registering an Oracle ACFS file system
 - [mount](#) (page 16-30) (Linux environments) or [acfsmountvol](#) (page 16-54) (Windows) for information about commands to mount an Oracle ACFS file system
 - *Oracle Database Reference* for information about the `V$ASM_VOLUME` view
 - *Oracle Real Application Clusters Administration and Deployment Guide* for information about Server Control Utility (SRVCTL)
-

16.1.3 Accessing an Oracle ACFS File System on a Different Node in the Cluster

If the node is part of a cluster, 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.

For example:

```
ASMCMD [+] > volenable -G data volumel
```

2. View information about the volume that you created on node 1.

For example:

```
ASMCMD [+] > volinfo -G data volumel
```

3. Mount the file system using the Oracle ACFS mount command.

For example:

```
# /bin/mount -t acfs /dev/asm/volumel-123 /acfsmounts/acfs1
```

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, 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 /acfsmounts/acfs1/myfile  
Oracle ACFS File System
```

The contents should match the file created previously on node 1.

See Also:

- [About Privileges for Oracle ASM](#) (page 3-28) for information about operating system groups and privileges.
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about the `volenable` command
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about the `volinfo` command
 - [mount](#) (page 16-30) (Linux environments) or [acfsmountvol](#) (page 16-54) (Windows) for information about commands to mount Oracle ACFS file systems
-
-

16.1.4 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 /acfsmounts/acfs1
```

See "[acfsutil snap create](#) (page 16-127)".

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" > /acfsmounts/acfs1/myfile
```

3. List the contents of the test file and the snapshot view of the test file.

For example:

```
$ cat /acfsmounts/acfs1/myfile
```

```
$ cat /acfsmounts/acfs1/.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.

16.1.5 Securing Oracle ACFS File Systems

The basic operations to manage security for an Oracle ACFS file system on Linux are discussed in this topic.

The scenario in this topic shows how to use Oracle ACFS security to ensure that only the maintenance user can access medical history files during the maintenance period. Oracle ACFS encryption is also enabled on the same file system.

In this scenario, the disk group on which the volume is created for the file system has compatibility attributes for ASM and ADVM set to 11.2.0.3 or higher.

For the examples in this section, various operating system users, operating system groups, and directories must exist.

The basic steps to manage security are:

1. Initialize security for Oracle ACFS.

Run the `acfsutil sec init` command to configure storage for security credentials and identify an operating system user as the first security administrator and the operating system security group. The security administrator must belong to the operating system group. This command must be run before any other security command and requires root or Windows Administrator privileges to run.

The `acfsutil sec init` command is only run once to set up Oracle ACFS security for each cluster and can be run from any node in the cluster. Other security commands can also be run from any node in a cluster. Security administrators are common for all Oracle ACFS file systems in a cluster.

For example, the following command initializes security for a cluster and creates the first security administrator (`medHistAdmin1`).

```
# /sbin/acfsutil sec init -u medHistAdmin1 -g medHistAdminGrp
```

The `medHistAdmin1` security administrator must belong to the `medHistAdminGrp` operating system group. That group is identified as the security group for the security administrators.

When the root user or Windows Administrator user runs the command, the user assigns a security password to the security administrator. The first security administrator can change the password with the `acfsutil sec admin password` command.

All `acfsutil sec` commands (other than `acfsutil sec init`) must be run by an Oracle ACFS security administrator and the administrator is prompted for the security administrator's password when each command is run.

Note:

When prompting for the security administrator's password, the following text displays: Realm management password

The password required is the Oracle ACFS security administrator's password, not the operating system password of the user.

Security administrators are allowed to browse all directories in an Oracle ACFS file system whether they have the underlying operating system permissions and whether any realm checks allow it. This functionality enables a security administrator to check the location of the files when securing them with Oracle ACFS security realms. However, a security administrator cannot view the contents of individual files without the appropriate operating system and security realm permissions.

2. Add additional security administrators as necessary.

The first security administrator can add additional security administrators to administer Oracle ACFS security with the `acfsutil sec admin add` command.

For example, add a new security administrator `medHistAdmin2`.

```
$ /sbin/acfsutil sec admin add medHistAdmin2
```

The `medHistAdmin2` user must belong to the operating system group (`medHistAdminGrp`) identified as the security administrator group with the `acfsutil sec init` command.

The `medHistAdmin2` security administrator should change the assigned temporary security password with the `acfsutil sec admin password` command. The `medHistAdmin2` administrator can add new security administrators.

3. Prepare an Oracle ACFS file system for security.

Run the `acfsutil sec prepare` on an Oracle ACFS file system before adding any security realms.

For example, prepare the Oracle ACFS file system mounted on `/acfsmounts/acfs1` for Oracle ACFS security.

```
$ /sbin/acfsutil sec prepare -m /acfsmounts/acfs1
```

By default, security is enabled for a file system after running this command. You can explicitly disable or enable security with the `acfsutil sec disable` or `acfsutil sec enable` commands.

This command automatically creates several security realms, such as the `SYSTEM_BackupOperators` security realm. Administrators can add users to the `SYSTEM_BackupOperators` realm which gives those users permissions to make backups of realm-secured files in the Oracle ACFS file system.

4. Provide encryption for this file system.

Encrypting the file system is optional, but is enabled in this scenario.

- a. First, run the `acfsutil encr init` command to initialize encryption and create the storage necessary for the encryption keys. This command must be run one time for each cluster on which encryption is set up.

For example, the following command initializes encryption for a cluster.

```
# /sbin/acfsutil encr init
```

This command must be run before any other encryption command and requires root or administrator privileges to run.

- b. Next, run the `acfsutil encr set` command to set encryption for the Oracle ACFS file system.

For example, the following command sets encryption for the file system mounted on the `/acfsmounts/acfs1` directory.

```
# /sbin/acfsutil encr set -m /acfsmounts/acfs1/
```

The `acfsutil encr set` command transparently generates a volume encryption key which is stored in the key store that was previously configured with the `acfsutil encr init` command. This command requires root or administrator privileges to run.

5. Create a security realm on the file system.

Run the `acfsutil sec realm create` command to create a security realm for a file system.

For example, create a security realm named `medHistRealm` which contains medical records files with all files encrypted in the realm.

```
$ /sbin/acfsutil sec realm create medHistRealm -m /acfsmounts/acfs1/
-e on -a AES -k 128
```

The `-e` option specifies that all the files in the realm are encrypted with the AES algorithm and the key length set to 128 bits. The file system must first be prepared for encryption with the `acfsutil encr init` and `acfsutil encr set` commands. You do not have to enter the same value for the `-k` option with `acfsutil sec realm create` as you have entered with the `acfsutil encr set` command.

6. Create security rules.

Run the `acfsutil sec rule create` command to create rules which determine access to the files and directories of a security realm.

For example, create rules that can enable the `medMaintenance` user to access medical records for the time period 10 PM to 2 AM for file maintenance. Also, create rules that can deny operations during the time period 8 AM to 9 AM and deny operations to the `medBrowse` user.

```
$ /sbin/acfsutil sec rule create medHistRule1a -m /acfsmounts/acfs1/
-t time 22:00:00,02:00:00 -o ALLOW
```

```
$ /sbin/acfsutil sec rule create medHistRule1b -m /acfsmounts/acfs1/
-t username medMaintenance -o ALLOW
```

```
$ /sbin/acfsutil sec rule create medHistRule1c -m /acfsmounts/acfs1/
-t time 08:00:00,09:00:00 -o DENY
```

```
$ /sbin/acfsutil sec rule create medHistRule1d -m /acfsmounts/acfs1/
-t username medBrowse -o DENY
```

You can edit rules with the `acfsutil sec rule edit` command.

7. Create security rule sets and add rules to rule sets.

Run the `acfsutil sec ruleset create` command to create rule sets to which rules can be added.

For example, create rule sets named `medRuleSet1` and `medRuleSet2` that include rules for operations on the files and directories of the security `medHistRealm` realm.

```
$ /sbin/acfsutil sec ruleset create medRuleSet1 -m /acfsmounts/acfs1/
```

```
$ /sbin/acfsutil sec ruleset create medRuleSet2 -m /acfsmounts/acfs1/
```

Add existing rules to the rule sets.

```
$ /sbin/acfsutil sec ruleset edit medRuleSet1 -m /acfsmounts/acfs1/
-a medHistRule1a,medHistRule1b -o ALL_TRUE
```

```
$ /sbin/acfsutil sec ruleset edit medRuleSet2 -m /acfsmounts/acfs1/
-a medHistRule1c,medHistRule1d -o ALL_TRUE
```

The `ALL_TRUE` option is the default action, but is added here to emphasize that both rules in each rule set must be true.

8. Add objects to a security realm.

Run the `acfsutil sec realm add` command to add objects, such as command rules, rule sets, and files, to a security realm. For example, add the `medRuleSet1` and `medRuleSet2` rule sets and all the files in the `/acfsmounts/acfs1/medicalrecords` directory to the `medHistRealm`.

When adding a rule set to a realm, the rule set is added with a command rule, such as `DELETEDFILE:medRuleSet1`. Only one rule set can be included with each command rule. To display a list of the command rules, use `acfsutil sec info` with the `-c` option.

The following `acfsutil sec realm add` command enables the `medMaintenance` user to delete medical records during the time period 10 PM to 2 AM, but blocks writing to files during 8 AM to 9 AM.

```
$ /sbin/acfsutil sec realm add medHistRealm -m /acfsmounts/acfs1/
-l DELETEDFILE:medRuleSet1
-f -r /acfsmounts/acfs1/medicalrecords
```

This `acfsutil sec realm add` command prevents the `medBrowse` user from writing or deleting medical records anytime.

```
$ /sbin/acfsutil sec realm add medHistRealm -m /acfsmounts/acfs1/
-l WRITE:medRuleSet2
-f -r /acfsmounts/acfs1/medicalrecords
```

This `acfsutil sec realm add` command adds backup operators to the `SYSTEM_BackupOperators` security realm that was automatically created with the `acfsutil sec prepare` command.

```
$ /sbin/acfsutil sec realm add SYSTEM_BackupOperators -m /acfsmounts/acfs1/
-G sysBackupGrp
```

Users that belong to the `sysBackupGrp` operating system group can now make backups of realm-secured files in the Oracle ACFS file system.

9. Display security information.

Run the `acfsutil sec info` command to display information for a security realm. For example, display security information for the `medHistRealm` realm.

```
$ /sbin/acfsutil sec info -m /acfsmounts/acfs1/ -n medHistRealm
```

To display the security realms to which a file or a directory belongs, run the `acfsutil sec info file` command. For example:

```
$ /sbin/acfsutil sec info file -m /acfsmounts/acfs1/  
/acfsmounts/acfs1/medicalrecords
```

10. Save security metadata as a backup.

Run the `acfsutil sec save` command to save the security metadata of a file system.

For example, save the security metadata of the `/acfsmounts/acfs1` file system to the `acfs1_backup.xml` file.

```
$ /sbin/acfsutil sec save -m /acfsmounts/acfs1  
-p acfs1_backup.xml
```

The `acfs1_backup.xml` security metadata backup file is saved in the `/acfsmounts/acfs1/.Security/backup/` directory. The saved XML file can be loaded with the `acfsutil sec load` command.

You can run some `acfsutil sec` commands in a batch file with the `acfsutil sec batch` command. For example, you could create a batch file that contains a group of `acfsutil sec rule` and `acfsutil sec ruleset` commands.

Auditing and diagnostic data for Oracle ACFS security is saved to log files.

See Also:

- [Oracle ACFS Security](#) (page 11-22) for more information about Oracle ACFS security, including security log files
 - [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - [acfsutil sec prepare](#) (page 16-96) and [acfsutil sec init](#) (page 16-93) for information about commands to set up security
 - [acfsutil sec disable](#) (page 16-90) and [acfsutil sec enable](#) (page 16-91) for information about commands to enable and disable security
 - [acfsutil sec admin add](#) (page 16-87) and [acfsutil sec admin password](#) (page 16-88) for information about commands to manage security administration
 - [acfsutil encr init](#) (page 16-119) and [acfsutil encr set](#) (page 16-123) for information about commands to manage encryption
 - [acfsutil sec realm create](#) (page 16-107), [acfsutil sec realm add](#) (page 16-98), and [acfsutil sec realm delete](#) (page 16-108) for information about commands to manage security realms
 - [acfsutil sec rule create](#) (page 16-110) and [acfsutil sec rule edit](#) (page 16-112) for information about commands to manage security rules
 - [acfsutil sec ruleset create](#) (page 16-114) and [acfsutil sec ruleset edit](#) (page 16-116) for information about commands to manage security rule sets
 - [acfsutil sec info](#) (page 16-91) and [acfsutil sec info file](#) (page 16-92) for information about commands to display security information
 - [acfsutil sec save](#) (page 16-116) and [acfsutil sec load](#) (page 16-95) for information about commands to load and save security metadata
 - [acfsutil sec batch](#) (page 16-89) for information about running `acfsutil sec` commands in a batch file
-

16.1.6 Encrypting Oracle ACFS File Systems

Basic operations to manage encryption on an Oracle ACFS file system on Linux are discussed in this topic.

The examples in this section show a scenario in which the medical history files are encrypted in an Oracle ACFS file system. The steps in this section assume Oracle ACFS security is not configured for the file system; however, you can use both Oracle ACFS security and encryption on the same file system. If you decide to use both security and encryption, then both encryption and security must be initialized for the cluster containing the file system. After security is initialized on the file system, then an Oracle ACFS security administrator runs `acfsutil sec` commands to provide encryption for the file system.

Because the `acfsutil encr set` and `acfsutil encr rekey -v` commands modify the encryption key store, you should back up the Oracle Cluster Registry (OCR) after

running these commands to ensure there is an OCR backup that contains all of the volume encryption keys (VEKs) for the file system.

The disk group on which the volume is created for the file system has compatibility attributes for ASM and ADVM set to 11.2.0.3 or higher.

For the examples in this section, various operating system users, operating system groups, and directories must exist.

The basic steps to manage encryption are:

1. Initialize encryption.

Run the `acfsutil encr init` command to initialize encryption and create the storage necessary for the encryption keys. This command must be run one time for each cluster on which encryption is set up.

For example, the following command initializes encryption for a cluster.

```
# /sbin/acfsutil encr init
```

This command must be run before any other encryption command and requires root or administrator privileges to run.

2. Set encryption parameters.

Run the `acfsutil encr set` command to set the encryption parameters for the entire Oracle ACFS file system.

For example, the following command sets the AES encryption algorithm and a file key length of 128 for a file system mounted on the `/acfsmounts/acfs1` directory.

```
# /sbin/acfsutil encr set -a AES -k 128 -m /acfsmounts/acfs1/
```

The `acfsutil encr set` command also transparently generates a volume encryption key which is stored in the key store that was previously configured with the `acfsutil encr init` command.

This command requires root or administrator privileges to run.

3. Enable encryption.

Run the `acfsutil encr on` command to enable encryption for directories and files.

For example, the following command enables encryption recursively on all files in the `/acfsmounts/acfs1/medicalrecords` directory.

```
# /sbin/acfsutil encr on -r /acfsmounts/acfs1/medicalrecords
-m /acfsmounts/acfs1/
```

For users that have appropriate permissions to access files in the `/acfsmounts/acfs1/medicalrecords` directory, they can still read the decrypted files.

This command can be run by an administrator or the file owner.

4. Display encryption information.

Run the `acfsutil encr info` command to display encryption information for directories and files.

```
# /sbin/acfsutil encr info -m /acfsmounts/acfs1/
-r /acfsmounts/acfs1/medicalrecords
```

This command can be run by an administrator or the file owner.

Auditing and diagnostic data for Oracle ACFS encryption is saved to log files. .

See Also:

- [Oracle ACFS Encryption](#) (page 11-25) for more information about Oracle ACFS encryption, including log files
 - [Securing Oracle ACFS File Systems](#) (page 16-9) for information about setting up security with encryption
 - [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - [acfsutil encr init](#) (page 16-119) for information about initializing encryption
 - [acfsutil encr set](#) (page 16-123) for information about setting encryption parameters
 - [acfsutil encr on](#) (page 16-121) for information about enabling encryption
 - [acfsutil encr info](#) (page 16-118) for information displaying encryption information
-
-

16.1.7 Tagging Oracle ACFS File Systems

The operations to manage tagging on directories and files in an Oracle ACFS file system on Linux are discussed in this topic.

The disk group on which the volume is created for the file system has compatibility attributes for ASM and ADVM set to 11.2.0.3 or higher.

Oracle ACFS implements tagging with Extended Attributes. There are some requirements when using Extended Attributes that should be reviewed.

The steps to manage tagging are:

1. Specify tag names for directories and files.

Run the `acfsutil tag set` command to set tags on directories or files. You can use these tags to specify which objects are replicated.

For example, add the `comedy` and `drama` tags to the files in the subdirectories of the `/acfsmounts/repl_data/films` directory.

```
$ /sbin/acfsutil tag set -r comedy /acfsmounts/repl_data/films/comedies
```

```
$ /sbin/acfsutil tag set -r drama /acfsmounts/repl_data/films/dramas
```

```
$ /sbin/acfsutil tag set -r drama /acfsmounts/repl_data/films/mysteries
```

In this example, the `drama` tag is purposely used twice and that tag is changed in a later step.

You must have system administrator privileges or be the file owner to run this command.

2. Display tagging information.

Run the `acfsutil tag info` command to display the tag names for directories or files in Oracle ACFS file systems. Files without tags are not displayed.

For example, display tagging information for files in the `/acfsmounts/repl_data/films` directory.

```
$ /sbin/acfsutil tag info -r /acfsmounts/repl_data/films
```

Display tagging information for files with the `drama` tag in the `/acfsmounts/repl_data/films` directory.

```
$ /sbin/acfsutil tag info -t drama -r /acfsmounts/repl_data/films
```

You must have system administrator privileges or be the file owner to run this command.

3. Remove and change tag names if necessary.

Run the `acfsutil tag unset` command to remove tags on directories or files. For example, unset the `drama` tag on the files in the `mysteries` subdirectory of the `/acfsmounts/repl_data/films` directory to apply a different tag to the subdirectory.

```
$ /sbin/acfsutil tag unset -r drama /acfsmounts/repl_data/films/mysteries
```

Add the `mystery` tag to the files in the `mysteries` subdirectory of the `/acfsmounts/repl_data/films` directory.

```
$ /sbin/acfsutil tag set -r mystery /acfsmounts/repl_data/films/mysteries
```

You must have system administrator privileges or be the file owner to run these commands.

See Also:

- [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - [Oracle ACFS Tagging](#) (page 11-39) for information about tagging an Oracle ACFS file system, including requirements for using Extended Attributes in tagging
 - [acfsutil tag set](#) (page 16-61) for information about specifying tag names
 - [acfsutil tag info](#) (page 16-60) for information about displaying tag name and details
 - [acfsutil tag unset](#) (page 16-63) for information about changing and removing tag names
-

16.1.8 Replicating Oracle ACFS File Systems

The operations to manage Oracle ACFS snapshot-based replication on an Oracle ACFS file system on Linux are discussed in this topic.

The disk groups on which volumes are created for the primary and standby file systems must have compatibility attributes for ASM and ADVM set to 12.2 or higher.

The steps to manage replication are:

1. Ensure that `ssh` has been configured for replication.
 - Ensure that host keys and user keys for the `ssh` command have been configured on your primary and standby clusters.
 - On Windows, ensure that Cygwin is installed and `ssh` is configured as required for replication.
2. Ensure that the snapshots needed by replication can be created at all times. At any given point replication, may need to be able to use two concurrent snapshots of the primary file system, and one snapshot of the standby file system.
3. Ensure that there is adequate network connectivity between the primary and standby sites. You should verify that the achievable network data transfer rate from primary to standby is substantially larger than the rate of change of data on the primary file system.

One way to estimate network data transfer rate is to start with an observed transfer rate, then reduce it to account for known sources of overhead. For example, you can calculate the elapsed time needed to FTP a 1 G file from the primary file system to the intended standby file system, during a period when network usage is low. This provides an estimate of the maximum achievable transfer rate. This rate should be reduced to account for overheads inherent in replication transfers, as well as to allow for other demands on the network. For replication overhead, a reasonable approach is to reduce the measured rate by 20%, then by an additional 5% for each node in the primary cluster.

To estimate the average rate of change on the primary, you can use the command `acfsutil info fs` with the `-s` option. This command should be run on each node where the primary file system is mounted, and displays that amount and rate of change to the file system on that node. To compute the total rate of change for the file system, the rate of change for each node must be aggregated. A reasonable value to use for `-s` is 900, which would yield a 15 minute sampling interval.

With the output from `acfsutil info fs` with the `-s` option, you can determine the average rate of change, the peak rate of change, and how long the peaks last. A conservative approach to using this data is to choose the peak rate of change as the target rate that must be accommodated.

Because replication must transfer all data changed on the primary to the standby, obviously the achievable network transfer rate must be higher, ideally significantly higher, than the target rate of change on the primary. If this is not the case, you should increase network capacity before implementing replication for this file system and workload.

For example, assume you have a four node primary cluster and you determine that a 1 G file can be transferred in 30 seconds, yielding a current FTP transfer rate of 33 M per second. An estimate of the current replication transfer rate would be approximately 20 M per second, calculated as follows:

$$33 \text{ MB/sec} * (1 - 0.2 - (4 * 0.05)) = 33 * 0.6 = \sim 20 \text{ MB/sec}$$

Also, you find that the average rate of change to the primary is 8 GB per hour, with a peak rate of 25 G per hour. Using the peak rate, you can calculate a target rate of change of approximately 7 M per second as follows:

$$(25 \text{ GB/hour} * 1024) / 3600 = \sim 7 \text{ MB/sec}$$

In the scenario that was discussed in this step, you can reasonably expect the network to be able to handle the additional workload from replication.

4. Ensure that there is adequate storage capacity on the primary and standby sites.

Estimate the storage capacity needed for replication on the sites hosting the primary and standby file systems. In the general case, the primary site must store two snapshots of the primary file system on an ongoing basis and the standby site must store a single snapshot of the standby file system. The space occupied by these snapshots mostly consists of user data or metadata preserved in the snapshot, that has since been modified which triggers a new copy of the data to be created.

The space occupied by replication-related snapshots can be directly viewed using the command `acfsutil snap info`. On the primary, check for snapshots with the names starting with the string `REPL`. On the standby, look for snapshots for names starting with `SDBACKUP`.

If you use interval-based replication, the `-i` option to `acfsutil repl init primary`, and if the replication operations are successfully completing within the specified interval, then the size of replication-related snapshots is related to the rate of change of the primary and the length of the interval. For example, with an average rate of change of 8 G per hour and a two hour replication interval, you would expect that snapshot storage usage is in the range of 16 G per snapshot.

Snapshot size does vary with the rate of change of the primary. Another factor is that snapshot size depends in part on the number of files in the file system, as well as the rate of change. Potentially more importantly, if you use constant mode replication, the `-C` option to `acfsutil repl init primary`, or if replication operations are not completing successfully in the interval given with interval-based replication because the interval is too small, the size of replication-related snapshots is difficult to predict in advance. In these cases, observe the size of the snapshots being generated over time and adjust the file system size as needed with the `acfsutil size` command to accommodate normal storage needs in addition in the presence of the snapshots. When collecting this information, a good starting point is to accommodate space for the snapshots to contain the data that is multiple times larger than the collection period, at the average rate of change of the primary.

While collecting this information, choose a conservative starting point for the amount of space to allow for replication snapshots. For example, you can compute the space needed to store changes to the file system over the collection period as described previously, then you can allocate several times that space for future snapshots.

5. Determine the user to be employed for replication and optionally set up tags.

Choose or create the replication user who logs in with `ssh` to the standby cluster to apply data replicated from the primary file system to the standby file system. This user is defined only at the OS level and not within Oracle. The user should belong to the groups defined for Oracle ASM administrator access.

Optionally set tags on directories and files to replicate only selected files in an Oracle ACFS file system. You can also add tags to files after replication has started.

6. Configure the site hosting the standby file system.

Before replicating an Oracle ACFS file system, configure the site hosting the standby file system by performing the following:

- Create a new *standby* file system of adequate size to hold the files replicated from the primary file system, as well as a single replication snapshot . For example:

```
/standby/repl_data
```

- Mount the file system on one node only.
- Run the `acfsutil repl init standby` command on the site hosting the standby file system. For example:

```
# /sbin/acfsutil repl init standby -u repluser /standby/repl_data
```

Note:

If the `acfsutil repl init standby` command is interrupted for any reason, the user must re-create the file system, mount it on one node only, and re-run the command.

This command requires the name of the replication user and the standby file system. The specified user is the user under which `ssh`, invoked from the primary cluster, logs in to the standby cluster to apply changes. This user is specified with the `-u` option. For example: `-u repluser`.

The mount point is the standby file system. For example: `/standby/repl_data`.

- After the `acfsutil repl init standby` command has successfully completed, mount the specified file system on all nodes of the standby cluster.

The `acfsutil repl init standby` command requires `root` or system administrator privileges to run.

7. After the standby file system has been set up, configure the site hosting the primary file system and start replication.

Run the `acfsutil repl init primary` command on the site hosting the primary file system. For example:

```
$ /sbin/acfsutil repl init primary -i 2h -s repluser@standby12_vip -m /standby/repl_data /acfsmounts/repl_data
```

This command requires the following configuration information.

- A replication interval, given with the option `-i interval` for interval mode or the option `-C` for constant mode replication. If an interval is specified, the option value is the minimum amount of time that elapses between replication operations. At the start of each operation, replication takes a new snapshot of the primary and compares it to the previous snapshot, if any. The changes needed to update the standby to match the primary are then sent to the standby. If `-C` is given instead of `-i interval`, a new replication operation is started as soon as the previous one completes.

For example, to set up a replication interval of two hours, specify `-i 2h`.

- The user name and network endpoint (VIP name or address, or host name or address) to be used to connect to the site hosting the standby file system, specified with the `-s` option. For example: `-s repluser@standby12_vip`

- The mount point of the primary file system. For example: `/acfsmounts/repl_data`
- If the mount point is different on the site hosting the standby file system than it is on the site hosting the primary file system, specify the mount point on the standby file system with the `-m standby_mount_point` option. For example: `-m /standby/repl_data`

The `acfsutil repl init primary` command requires root or system administrator privileges to run.

8. Monitor information about replication on the file system.

The `acfsutil repl info` command displays information about the state of the replication processing on the primary or standby file system.

For example, you can run the following on the site hosting the primary file system to display configuration information.

```
$ /sbin/acfsutil repl info -c -v /acfsmounts/repl_data
```

You must have system administrator (the user root on non-Windows systems or local SYSTEM on Windows) or Oracle ASM administrator privileges to run this command.

9. Manage the replication background process.

Run the `acfsutil repl bg` command to start, stop, or retrieve information about the replication background process.

For example, run the following command to display information about the replication process for the `/acfsmounts/repl_data` file system.

```
$ /sbin/acfsutil repl bg info /acfsmounts/repl_data
```

You must have system administrator or Oracle ASM administrator privileges to run the `acfsutil repl bg info` command.

10. Pause replication momentarily if necessary.

Run the `acfsutil repl pause` to momentarily stop replication. Run the `acfsutil repl resume` command as soon as possible to resume replication.

For example, the following command pauses replication on the `/acfsmounts/repl_data` file system.

```
$ /sbin/acfsutil repl pause /acfsmounts/repl_data
```

The following command resumes replication on the `/acfsmounts/repl_data` file system.

```
$ /sbin/acfsutil repl resume /acfsmounts/repl_data
```

You must have system administrator or Oracle ASM administrator privileges to run the `acfsutil repl pause` and `acfsutil repl resume` commands.

11. Failing over to a standby or turning a standby file system into an active file system.

If the primary file system is inaccessible, you can run `acfsutil repl terminate standby mount_point` to turn the standby file system into an active file system. If the primary file system still exists, you should terminate the primary first with `acfsutil repl terminate primary mount_point`.

Before terminating replication with `acfsutil repl terminate standby` on the standby file system, you can determine the point in time of the primary file system that the standby file system represents. This timestamp is displayed with `acfsutil repl info -c` as `Last sync time with primary`. If the failover action must be coordinated with Oracle Data Guard, you can use the timestamp to set back the database if needed, or perform other necessary actions that are based on the timestamp.

The standby file system may be in the process of being modified by replication. This could occur if:

- The primary file system is available and a replication operation is currently in progress.
- The primary file system is not available, but a replication operation was in progress when it became unavailable.

To be sure of obtaining the contents of the standby file system when it was last identical to a snapshot from the primary, follow one of these procedures.

- If the primary file system is available, run the `acfsutil repl terminate primary` command on the primary site to terminate replication. The command waits for any in-progress replication operation to complete before it returns. Then run `acfsutil repl info -c` to determine the point in time of the primary file system represented on the standby. After you have this information, run `acfsutil repl terminate standby` on the standby site.
- If the primary file system is not available, you should first compare two date strings from the `acfsutil repl info -c` output; these are the dates from the `Receiving primary as of line` and the `Last sync time with primary line`. If these dates are identical, then the standby file system contains the most recent available point-in-time image of the primary. If they are not equal, you must use the backup snapshot recorded by replication to recover the last point-in-time image captured on the standby. You can find this snapshot using the `acfsutil snap info` command. Search for a snapshot with a name of the form:

```
SDBACKUP_tstamp1_REPL_tstamp2
```

where `tstamp1` represents the time at which the backup snapshot was created, and `tstamp2` represents the point in time when the primary contents in this snapshot were recorded. There should be only one backup snapshot present. The date of the backup snapshot corresponds to the date in the `Last sync time with primary` line output by `acfsutil repl info -c`. If there is no backup snapshot available, then the contents of the primary were never successfully transferred to the standby.

The date of the backup snapshot corresponds to the date in the `Last sync time with primary` line output by `acfsutil repl info -c`.

To use the backup snapshot, you must terminate replication and ensure that the snapshot is preserved because it is deleted by default. To ensure the snapshot is preserved, add the `-k` option to the command line for `acfsutil repl terminate standby`. After replication has been terminated, you can run the `acfsutil snap remaster` command to use the snapshot as the new contents of the standby file system.

Note:

When replication is in use, replication snapshots can be viewed using the `acfsutil snap info` command, just as any other snapshot can. You can use this command to get an approximate idea of the space currently occupied by replication snapshots.

See Also:

- [Oracle ACFS Replication](#) (page 11-36) for information about replicating an Oracle ACFS file system
 - [Disk Group Compatibility](#) (page 4-50) for information about disk group compatibility
 - [Configuring ssh for Use With Oracle ACFS Replication](#) (page 18-1) and [Installing ssh and Cygwin on Windows](#) (page 18-5) for information about configuring replication
 - [acfsutil repl init](#) (page 16-69) for information initiating replication
 - [About Privileges for Oracle ASM](#) (page 3-28) for information about user privileges for Oracle ASM
 - [acfsutil repl bg](#) (page 16-64) for more information about managing replication background operations
 - [acfsutil repl pause](#) (page 16-72) and [acfsutil repl resume](#) (page 16-73) for more information about pausing and resuming replication operations
 - [acfsutil repl info](#) (page 16-67) for information about displaying replication details
 - [Creating an Oracle ACFS File System](#) (page 16-5) for information about creating a file system
 - [Tagging Oracle ACFS File Systems](#) (page 16-16) for information about the steps to tag files
-
-

16.1.9 Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems

This section 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](#) (page 16-23)
- [Dismounting an Oracle ACFS File System](#) (page 16-24)
- [Disabling a Volume](#) (page 16-24)

16.1.9.1 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 /acfsmounts/acfs1
```

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](#) (page 11-11)". For information about `acfsutil registry`, see "[acfsutil registry](#) (page 16-160)".

16.1.9.2 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 /acfsmounts/acfs1

# /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](#) (page 16-32)" or "[acfsdismount](#) (page 16-52)".

Use `fsck` on Linux systems or `acfschkdisk` on Windows systems to check a file system. For information about the commands to check a file system, see "[fsck](#) (page 16-26)" or "[acfschkdisk](#) (page 16-50)".

16.1.9.3 Disabling a Volume

To disable a volume, you must first dismount the file system on which the volume is mounted.

For example:

```
# /bin/umount /acfsmounts/acfs1
```

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 [Managing Oracle ADVM with ASMCMD](#) (page 15-1).

16.1.10 Removing an Oracle ACFS File System and a Volume

You can remove an Oracle ACFS file system and volume with `acfsutil` and `ASMCMD` commands.

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 /acfsmounts/acfs1
acfsutil registry: successfully removed ACFS mount point
/acfsmounts/acfs1 from Oracle Registry
```

2. Dismount the file system.

For example:

```
# /bin/umount /acfsmounts/acfs1
```

You must dismount the file system on all nodes of a cluster.

Use `umount` on Linux systems or `acfsdismount` on Windows systems.

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
```

4. Optionally you can disable the volume with the `ASMCMD voldisable` command.

For example:

```
ASMCMD> voldisable -G data volume1
```

5. Delete the volume with the `ASMCMD voldelete` command.

For example:

```
ASMCMD> voldelete -G data volume1
```

See Also:

- [acfsutil registry](#) (page 16-160) for information about running `acfsutil registry`
 - [umount](#) (page 16-32) or [acfsdismount](#) (page 16-52) for information about running the `umount` or `acfsdismount` commands
 - [acfsutil rmfs](#) (page 16-162) for information about running the `acfsutil rmfs` command
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about running the `voldisable` command
 - [Managing Oracle ADVM with ASMCMD](#) (page 15-1) for information about running the `voldelete` command
-

16.2 Oracle ACFS Command-Line Tools for Linux Environments

This topic provides a summary of the Oracle ACFS commands for Linux environments.

[Table 16-1](#) (page 16-26) lists the Oracle ACFS commands for Linux environments with brief descriptions. The commands in [Table 16-1](#) (page 16-26) have been extended with additional options to support Oracle ACFS. All other Linux 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 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 16-1](#) (page 16-26) 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 16-1](#) (page 16-26).

Note:

When using Security-Enhanced Linux (SELinux) in enforcing mode with Oracle ACFS, ensure that the Oracle ACFS file systems are mounted with an SELinux default context. Refer to your Linux vendor documentation for information about the context mount option.

Table 16-1 Summary of Oracle ACFS commands for Linux environments

Command	Description
fsck (page 16-26)	Checks and repairs an Oracle ACFS file system.
mkfs (page 16-28)	Creates an Oracle ACFS file system.
mount (page 16-30)	Mounts an Oracle ACFS file system.
umount (page 16-32)	Dismounts an Oracle ACFS file system.

16.2.1 fsck

Purpose

Checks and repairs an Oracle ACFS file system.

Syntax and Description

```
fsck -t acfs -h /dev/null
fsck [-a|-f] [-v] -t acfs [-n|-y] [-x file_name] volume_device
```

`fsck -t acfs -h /dev/null` displays usage text and exits.

[Table 16-2](#) (page 16-26) contains the options available with the `fsck` command.

Table 16-2 Options for the Linux fsck command

Option	Description
<code>-a</code>	Specifies to automatically repair the file system.

Table 16-2 (Cont.) Options for the Linux fsck command

Option	Description
-f	Forces the file system into mountable state without completing a file system check or repair.
-v	Specifies verbose mode. The progress is displayed as the operation occurs. Running in verbose mode can impact performance.
-h	Displays the usage help text and exits.
-n	Answers no to any prompts.
-y	Answers yes to any prompts.
-x <i>file_name</i>	Specified for accelerator data collected by <code>acfsutil meta</code> . Only used for this type of data.
<i>volume_device</i>	Specifies the primary Oracle ADVM volume device.

`fsck` checks and repairs an existing Oracle ACFS. This command can only be run on a dismounted file system. `root` privileges are required to run `fsck`. The Oracle ACFS driver must be loaded for `fsck` to work.

By default, `fsck` only checks for and reports any errors. The `-a` flag must be specified to instruct `fsck` to repair errors in the file system. Do not interrupt `fsck` during the repair operation.

In a few cases, `fsck` prompts for questions before proceeding to check a file system. These cases include:

- If `fsck` detects that another `fsck` is in progress on the file system
- If `fsck` detects that the Oracle ACFS driver is not loaded
- If the file system does not appear to be Oracle ACFS

In checking mode, `fsck` also prompts if there are transaction logs that have not been processed completely due to an incomplete shutdown. To run in a non-interactive mode, include either the `-y` or `-n` options to answer yes or no to any questions.

`fsck` 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, `fsck` attempts to write to the current working directory. The files that `fsck` creates are roughly the size of the file system being checked divided by 32K. At most three such files are allocated. For example, a 2 GB file system being checked causes `fsck` to generate one to three 64K working files in the `/usr/tmp` directory. These files are deleted after `fsck` has finished.

In the event that `fsck` 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 `fsck` is run in `fix` mode. For security reasons only the `root` user on Linux can 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 file](#) (page 16-147)".

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 16-2 Using the `fsck` command

```
# /sbin/fsck -a -y -t acfs /dev/asm/volume1-123
```

16.2.2 mkfs

Purpose

Creates an Oracle ACFS file system.

Syntax and Description

```
mkfs -t acfs -h
mkfs [-v] [-f] -t acfs [-i {512 | 4096}] [-n name ] [-a accelerator_volume] [-c
release_version]
      volume_device [size]
```

`mkfs -t acfs -h` displays usage text and exits.

[Table 16-3](#) (page 16-28) contains the options available with the `mkfs` command.

Table 16-3 Options for the Linux `mkfs` 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>-i {512 4096}</code>	Specifies a file system with metadata block size of 512 or 4096 bytes.
<code>-n name</code>	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.

Table 16-3 (Cont.) Options for the Linux mkfs command

Option	Description
-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. Use this option with caution.
-h	Displays the usage help text and exits.
<i>size</i>	Specifies the size that the file system should consume on the named device. The quantity specified can be in units of K (kilobytes), M (megabytes), G (gigabytes), T (terabytes), or P (petabytes). If a unit is not specified, the default is bytes. If this option is not specified, the entire device is consumed.
-a <i>accelerator_volume</i>	Specifies a secondary Oracle ACFS accelerator volume that is used by Oracle ACFS to store critical metadata.
-c <i>release_version</i>	Creates an Oracle ACFS file system with Oracle ACFS release compatibility set to the value of <i>release_version</i> .
<i>volume_device</i>	Specifies an existing Oracle ADVM device file that is to be formatted. This device is the primary volume.

You can use `mkfs` to create the on disk structure needed for Oracle ACFS file system to be mounted. The `mkfs` command is the traditional Linux command used to build a file system. After `mkfs` runs successfully, the `USAGE` column in the `V$ASM_VOLUME` view displays `ACFS`. Oracle ADVM volumes are created with the `ASMCMD volcreate` command. For information about the `volcreate` command, refer to "[volcreate](#) (page 15-2)".

The value specified with the `-c release_version` option must be greater than or equal to the value of `COMPATIBLE.ADVM` for the disk group, and must be less than or equal to the running Oracle Grid infrastructure release version. If `-c release_version` is not specified, the value of `COMPATIBLE.ADVM` is used. After the compatibility is set, it cannot be downgraded. The `-c release_version` option can be used in situations where it is not possible, or desirable, to update `COMPATIBLE.ADVM` and `COMPATIBLE.ASM` for the disk group, but you want to use an Oracle ACFS feature that requires a compatibility increase. After the Oracle ACFS compatibility is updated, it is no longer possible to mount the file system with an older Oracle Grid infrastructure release. To change the compatibility of an existing file system, refer to "[acfsutil compat set](#) (page 16-145)". For information about Oracle ASM disk group compatibility attributes, see "[Disk Group Compatibility](#) (page 4-50)".

The minimum file system size is 200 MB for the 512 byte metadata block size format and 512 MB for the 4 KB metadata block size format.

The Oracle ACFS driver must be loaded for `mkfs` to operate successfully.

`root` privilege is not required. The ownership of the volume device file dictates who can run this command.

Oracle ACFS Accelerator Volume

Using an accelerator volume can improve performance by reducing the time to access and update Oracle ACFS metadata. You should create the accelerator volume on a

disk group with storage that is significantly faster than the storage of the primary volume. The accelerator volume is created with the `-a` option of the `mkfs` command. For more information, refer to [Oracle ACFS Accelerator Volume](#) (page 11-43).

4 KB Sector/Metadata

The following list summarizes the use of the `-i` option and metadata block size based on the value of the `COMPATIBLE.ADVM` disk group attribute.

- If `COMPATIBLE.ADVM` is set to `12.2` or greater, then the metadata block size is 4096 bytes by default.
- If `COMPATIBLE.ADVM` is set to less than `12.2`, then the block size is set to 512 bytes.
- If `COMPATIBLE.ADVM` is set to less than `12.2` and the logical sector size is not 512 bytes, then the command fails.

User data IO continues to support transfers as small as 512 bytes for normal user IO requests. When the Oracle ADVM volume of the file system has a logical disk sector size of 4 K, user Direct IO requests should be aligned on 4 K file offsets and be a multiple of 4 KB lengths for the best performance. Note that only 4 K metadata block size formats support ADVM volumes with 4 K logical disk sectors.

The `COMPATIBLE.ADVM` value must be set to `12.2` or greater to format an Oracle ADVM volume whose logical disk size is 4096 or to use the `-i 4096` option to format an Oracle ADVM volume whose logical disk sector size is 512 bytes.

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. For example:

```
ASMCMD [+] > volinfo -a
...
      Volume Name: VOLUME1
      Volume Device: /dev/asm/volume1-123
      State: ENABLED
...
```

For information about the `volcreate` and `volinfo` commands, refer to [Managing Oracle ADVM with ASMCMD](#) (page 15-1).

This command creates an Oracle ACFS file system on a volume device file.

Example 16-3 Using the `mkfs` command

```
$ /sbin/mkfs -t acfs /dev/asm/volume1-123
```

Example 16-4 Using the `mkfs` command to create an accelerator volume

This command creates an Oracle ACFS file system with an accelerator volume.

```
$ /sbin/mkfs -t acfs -a /dev/asm/volume2-130 /dev/asm/volume1-127
```

16.2.3 mount

Purpose

Mounts an Oracle ACFS file system.

Syntax and Description

```
mount -h
mount [-v] -n -t acfs [-o options] volume_device mount_point
mount
```

`mount -h` displays usage text and exits.

Table 16-4 (page 16-31) contains the options available with the `mount` command.

Table 16-4 Options for the Linux `mount` 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>-n</code>	Specifies not to update the <code>/etc/mtab</code> file.
<code>-o</code>	Options are specified with the <code>-o</code> flag followed by a comma-delimited string of options. The following options are available: <ul style="list-style-type: none"> <code>all</code> Reads the Oracle ACFS managed file systems in the Oracle Clusterware namespace, mounting all file systems with the <code>AUTO_START=always</code>. When the <code>-o all</code> option is specified, other <code>-o</code> options are ignored. 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>nodeleteopen</code> Fails the deletion of any file that is open.
<code>volume_device</code>	Specifies a primary 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>mount_point</code>	Specifies the directory where the file system is mounted. This directory must exist before you run the <code>mount</code> command.

`mount` attaches a file system to the Oracle ACFS hierarchy at the mount point that is the 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` runs successfully, the `MOUNTPATH` field in the `V$ASM_VOLUME` view displays the directory name on which the file system is now mounted.

An Oracle ACFS file system should only be mounted on one mount point. The same mount point name should be used on all cluster members.

The `mount` command lists all mounted file systems if it is run with no parameters.

`root` privilege is required to run `mount`.

Examples

The first example shows how to mount `volume1-123` on the mount point `/acfsmounts/acfs1`. 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 16-5 Using the mount command

```
# /bin/mount -t acfs /dev/asm/volume1-123 /acfsmounts/acfs1

# /bin/mount -t acfs -o all none none
```

16.2.4 umount

Purpose

Dismounts an Oracle ACFS file system.

Syntax and Description

```
umount -h
umount [-v] -l -n volume_device |mount_point
umount -a -l -n [-t acfs]
```

`umount -h` displays usage text and exits.

[Table 16-5](#) (page 16-32) contains the options available with the `umount` command.

Table 16-5 Options for the Linux 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.

Table 16-5 (Cont.) Options for the Linux umount command

Option	Description
-l	Specifies a lazy unmount. The Oracle ACFS file system is dismounted, but the cleanup of all references to the file system occurs when the system is no longer busy.
-n	Specifies not to update the <code>/etc/mntab</code> file.
<code>volume_device</code>	Specifies an Oracle ADVM volume device file that has been formatted by <code>mkfs</code> .
<code>mount_point</code>	Specifies the directory where the file system is mounted. This directory must exist before you run the <code>mount</code> command.

`umount` detaches an Oracle ACFS from the file system hierarchy on the current node.

The `umount` command does not return control back to the operating system prompt until the file system has been completely unmounted.

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 16-6 Using the umount command

```
# /bin/umount /dev/asm/volume1-123

# /bin/umount /acfsmounts/acfs1
```

16.3 Oracle ACFS Command-Line Tools for the Solaris Environment

This topic provides a summary of the Oracle ACFS commands for Solaris.

Note:

If a disk group is dismounted with the force option on Solaris, any Oracle ADVM volumes device files that were on that disk group remain on the system. These files are removed when the disk group is remounted.

[Table 16-6](#) (page 16-33) lists the Oracle ACFS commands for Solaris with brief descriptions. The commands in [Table 16-6](#) (page 16-33) have been extended with additional options to support Oracle ACFS on Solaris.

Table 16-6 Summary of Oracle ACFS commands for Solaris

Command	Description
fsck (page 16-34)	Checks and repairs an Oracle ACFS file system on Solaris.
mkfs (page 16-36)	Creates an Oracle ACFS file system on Solaris.

Table 16-6 (Cont.) Summary of Oracle ACFS commands for Solaris

Command	Description
mount (page 16-38)	Mounts an Oracle ACFS file system on Solaris.
umount/umountall (page 16-40)	Dismounts an Oracle ACFS file system on Solaris.

16.3.1 fsck

Purpose

Checks and repairs an Oracle ACFS file system on the Solaris operating system.

Syntax and Description

```
fsck -F acfs -o h /dev/null
fsck -F acfs [{-n|N}|{-y|Y}] [-o options] volume_device
```

`fsck -F acfs -o h /dev/null` displays usage text and exits.

[Table 16-7](#) (page 16-34) contains the options available with the `fsck` command.

Table 16-7 Options for the Solaris fsck command

Option	Description
<code>-F acfs</code>	Specifies the type of file system on Solaris. <code>acfs</code> designates the Oracle ACFS type.
<code>-n N</code>	Answers no to any prompts.
<code>-y Y</code>	Answers yes to any prompts.
<code>-o</code>	Specifies that options follow (<code>a</code> , <code>f</code> , <code>h</code> , <code>v</code>). Options are preceded with the <code>-o</code> flag and entered as a comma-delimited string. For example: <code>-o a, v</code> <ul style="list-style-type: none"> • <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>h</code> Displays the usage text and exits. • <code>v</code> Specifies verbose mode. The progress is displayed as the operation occurs. Running in verbose mode can impact performance.
<code>volume_device</code>	Specifies an Oracle ADVM device file.

`fsck` checks and repairs an existing Oracle ACFS file system. This command can only be run on a dismounted file system. `root` privileges are required to run `fsck`. The Oracle ACFS driver must be loaded for `fsck` to work.

By default, `fsck` only checks for and reports any errors. The `-o a` option must be specified to instruct `fsck` to fix errors in the file system. Do not interrupt `fsck` during the repair operation.

In a few cases, `fsck` prompts for questions before proceeding to check a file system. These cases include:

- If `fsck` detects that another `fsck` is in progress on the file system
- If `fsck` detects that the Oracle ACFS driver is not loaded
- If the file system does not appear to be Oracle ACFS

In checking mode, `fsck` also prompts if there are transaction logs that have not been processed completely due to an incomplete shutdown. To run in a non-interactive mode, include either the `-y` or `-n` options to answer yes or no to any questions.

`fsck` 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, `fsck` attempts to write to the current working directory. The files that `fsck` creates are roughly the size of the file system being checked divided by 32K. At most three such files are allocated. For example, a 2 GB file system being checked causes `fsck` to generate one to three 64K working files in the `/usr/tmp` directory. These files are deleted after `fsck` has finished.

In the event that `fsck` 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 `fsck` is run in fix mode. For security reasons only the `root` user on Linux can 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 file](#) (page 16-147)".

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 16-7 Using the `fsck` command

```
# /usr/sbin/fsck -F acfs -y -o a /dev/asm/volume1-123
```

16.3.2 mkfs

Purpose

Creates an Oracle ACFS file system on the Solaris operating system.

Syntax and Description

```
mkfs -F acfs -o h /dev/null
mkfs -F acfs [-o options] [-c release_version] volume_device [size]
```

`mkfs -F acfs -o h /dev/null` displays usage text and exits.

[Table 16-8](#) (page 16-36) contains the options available with the `mkfs` command.

Table 16-8 Options for the Solaris `mkfs` command

Option	Description
<code>-F acfs</code>	Specifies the type of file system on Solaris. <code>acfs</code> designates the Oracle ACFS type.
<code>-o options</code>	Specifies that options follow (<code>f</code> , <code>h</code> , <code>n</code> , <code>v</code>). Options are preceded with the <code>-o</code> flag and entered as a comma-delimited string. For example: <code>-o f,v</code> <ul style="list-style-type: none"> • <code>h</code> Displays the usage text and exits. • <code>f</code> 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. Use this option with caution. • <code>i=512 i=4096</code> Specifies a file system with metadata block size of 512 or 4096 bytes. • <code>n=name</code> 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. • <code>v</code> Specifies verbose mode. The progress is displayed as the operation occurs.
<code>-c release_version</code>	Creates an Oracle ACFS file system with Oracle ACFS release compatibility set to the value of <code>release_version</code> .
<code>volume_device</code>	Specifies an Oracle ADVM device file.
<code>size</code>	Specifies the size of the file system in 512-byte units or in units of <code>K</code> <code>M</code> <code>G</code> <code>T</code> <code>P</code> . Units specified are in <code>K</code> (kilobytes), <code>M</code> (megabytes), <code>G</code> (gigabytes), <code>T</code> (terabytes), or <code>P</code> (petabytes).

You can use `mkfs` to create the on disk structure needed for Oracle ACFS file system to be mounted. The `mkfs` command is the traditional Linux command used to build a file system. After `mkfs` runs 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 value specified with the `-c release_version` option must be greater than or equal to the value of `COMPATIBLE.ADVM` for the disk group, and must be less than or equal to the running Oracle Grid infrastructure release version. If `-c release_version` is not specified, the value of `COMPATIBLE.ADVM` is used. After the compatibility is set, it cannot be downgraded. The `-c release_version` option can be used in situations where it is not possible, or desirable, to update `COMPATIBLE.ADVM` and `COMPATIBLE.ASM` for the disk group, but you want to use an Oracle ACFS feature that requires a compatibility increase. After the Oracle ACFS compatibility is updated, it is no longer possible to mount the file system with an older Oracle Grid infrastructure release. To change the compatibility of an existing file system, refer to "[acfsutil compat set](#) (page 16-145)". For information about Oracle ASM disk group compatibility attributes, see "[Disk Group Compatibility](#) (page 4-50)".

The minimum file system size is 200 MB for the 512 byte metadata block size format and 512 MB for the 4 KB metadata block size format.

The Oracle ACFS driver must be loaded for `mkfs` to work.

4 KB Sector/Metadata

The following list summarizes the use of the `-o i` option and metadata block size based on the value of the `COMPATIBLE.ADVM` disk group attribute.

- If `COMPATIBLE.ADVM` is set to `12.2` or greater, then the metadata block size is 4096 bytes by default.
- If `COMPATIBLE.ADVM` is set to less than `12.2`, then the block size is set to 512 bytes.
- If `COMPATIBLE.ADVM` is set to less than `12.2` and the logical sector size is not 512 bytes, then the command fails.

User data IO continues to support transfers as small as 512 bytes for normal user IO requests. When the Oracle ADVM volume of the file system has a logical disk sector size of 4 K, user Direct IO requests should be aligned on 4 K file offsets and be a multiple of 4 KB lengths for the best performance. Note that only 4 K metadata block size formats support ADVM volumes with 4 K logical disk sectors.

The `COMPATIBLE.ADVM` value must be set to `12.2` or greater to format an Oracle ADVM volume whose logical disk size is 4096 or to use the `-o i=4096` option to format an Oracle ADVM volume whose logical disk sector size is 512 bytes.

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 [Managing Oracle ADVM with ASMCMD](#) (page 15-1).

Next create an Oracle ACFS file system on the volume device file.

Example 16-8 Using the mkfs command

```
$ /usr/sbin/mkfs -F acfs /dev/asm/volume1-123
```

16.3.3 mount

Purpose

Mounts an Oracle ACFS file system on the Solaris operating system.

Syntax and Description

```
mount -F acfs -o h /tmp /dev/null
mount -F acfs [-r] [-o options] volume_device mount_point
```

`mount -F acfs -o h /tmp /dev/null` displays usage text and exits.

[Table 16-9](#) (page 16-38) contains the options available with the `mount` command.

Table 16-9 Options for the Solaris mount command

Option	Description
<code>-F acfs</code>	Specifies the type of file system on Solaris. <code>acfs</code> designates the Oracle ACFS type.
<code>-r</code>	Mounts in read-only mode.

Table 16-9 (Cont.) Options for the Solaris mount command

Option	Description
-o	<p>Specifies that options follow. Options are preceded with the <code>-o</code> flag followed by a comma-delimited string of options. For example: <code>-o all,v</code></p> <p>The following options are available:</p> <ul style="list-style-type: none"> • <code>all</code> Reads the Oracle ACFS managed file systems in the Oracle Clusterware namespace, mounting all file systems with the <code>AUTO_START=always</code>. The <code>-o all</code> option requires two placeholders to satisfy the command arguments: a dummy argument for the volume device and any valid directory. When the <code>-o all</code> option is specified, other <code>-o</code> options are ignored. 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>devices/nodevices</code> Allows or disallows the opening of any character or block special devices from this mount. • <code>exec/noexec</code> Allows or disallows the execution of programs in this file system. • <code>h</code> Displays the usage help text and exits. • <code>rootsuid/norootsuid</code> Allows or disallows the execution of <code>setuid</code> to root executables (binaries by non-root users whose permissions allow set user Id execution, and are owned by root). <code>rootsuid</code> is the default action. If <code>norootsuid</code> is specified, an attempt to run these executables as a non-root user fails with a permission denied error. • <code>ro</code> Mounts the file system in read-only mode. • <code>rw</code> Mounts the file system in read-write mode. • <code>setuid/nosetuid</code> Allows or disallows the execution of <code>setuid</code> and <code>setgid</code> programs. • <code>suid/nosuid</code> Allows or disallows the execution of <code>setuid</code> and <code>setgid</code> programs and the opening of any character or block special devices from this mount. • <code>v</code> Specifies verbose mode. The progress is displayed as the operation occurs.
<i>volume_device</i>	<p>Specifies an Oracle ADVM volume device file that has been formatted by <code>mkfs</code>. <i>device</i> is required but can be a dummy value.</p>

Table 16-9 (Cont.) Options for the Solaris mount command

Option	Description
<code>mount_point</code>	Specifies the directory where the file system is mounted. This directory must exist before you run the <code>mount</code> command.

`mount` attaches a file system to the Oracle ACFS hierarchy at the mount point that is the name of a directory. The mount occurs 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` runs successfully, the `MOUNTPATH` field in the `V$ASM_VOLUME` view displays the directory name on which the file system is now mounted.

An Oracle ACFS file system should only be mounted on one mount point. The same mount point name should be used on all cluster members.

`root` privilege is required to run `mount`.

Examples

The first example shows how to mount `volume1-123` on the mount point `/acfsmounts/acfs1`. The second example shows how to mount all the registered Oracle ACFS file systems. Placeholder arguments must be provided for the volume device name and mount point when specifying the `-o all` option. The volume device can be a dummy value, such as `none`. Any valid directory can be specified for the mount point, such as `/tmp`.

Example 16-9 Using the mount command

```
# /sbin/mount -F acfs /dev/asm/volume1-123 /acfsmounts/acfs1

# /sbin/mount -F acfs -o all none /tmp
```

16.3.4 umount/umountall

Purpose

Dismounts an Oracle ACFS file system on the Solaris operating system.

Syntax and Description

```
umount -V [mount_point | volume_device]
umountall -F acfs
```

`umountall -F acfs` dismounts all Oracle ACFS file systems.

[Table 16-10](#) (page 16-41) contains the options available with the `umount` command.

Table 16-10 Options for the Solaris umount command

Option	Description
-v	Echoes the complete command line, but does not run the command. Use this option to verify and validate the command line before execution. Valid only with <code>umount</code> .
<code>mount_point</code>	Specifies the directory where the file system is mounted. Valid only with <code>umount</code> .
<code>volume_device</code>	Specifies the Oracle ADVM volume device name associated with the file system. Valid only with <code>umount</code> .
-F <code>acfs</code>	Specifies the type of file system on Solaris. <code>acfs</code> designates the Oracle ACFS type. Valid only with <code>umountall</code> .

`umount` and `umountall` detach an Oracle ACFS from the file system hierarchy on the current node.

The `umount` and `umountall` commands do not return control back to the operating system prompt until the file system has been completely unmounted.

If a file system is busy, `umount` and `umountall` fail.

root privileges are required to run the `umount` and `umountall` commands.

Examples

The following examples show how to dismount an Oracle ACFS file system. The first example specifies the mount point of the file system to dismount. The second example specifies the volume device associated with the file system to dismount. The third example dismounts all Oracle ACFS file systems.

Example 16-10 Using the umount command

```
# /sbin/umount /dev/asm/volume1-123

# /sbin/umount /acfsmounts/acfs1

# /sbin/umountall -F acfs
```

16.4 Oracle ACFS Command-Line Tools for the AIX Environment

This topic provides a summary of the Oracle ACFS commands for the AIX operating system.

During the installation of Oracle ACFS and Oracle ADVM with Oracle Grid Infrastructure on AIX, several system security classes (authorizations) are created for Oracle ACFS and Oracle ADVM operations. These authorizations enable the following operations:

- Configuration of Oracle ACFS and Oracle ADVM devices
- Removal of Oracle ACFS and Oracle ADVM devices
- Definition of Oracle ACFS and Oracle ADVM devices

These authorizations are collected in a role (`oracle_devmgmt`) which is assigned by default to the Oracle Grid Infrastructure user and Oracle ASM administrator. You can run the `lsrole` or `roledist` command after installation to see this new role. These

commands can be run at any time to ensure that the user has the proper authorization to manage Oracle ACFS and Oracle ADVM.

For example, run the following as the Oracle ASM administrator:

```
$ rolelist -e
  oracle_devmgmt
```

After the initial installation on an AIX operating system, the shell from which the root script was run does not inherit the new role. If the role is not present for the user, then the `swrole` command must be run before performing Oracle ACFS or Oracle ADVM operations.

For example, run the following as the Oracle ASM administrator:

```
$ swrole oracle_devmgmt
```

Alternatively, you can open a new shell to perform Oracle ACFS or Oracle ADVM operations.

During the removal of Oracle Grid Infrastructure, the `oracle_devmgmt` role and its associated authorizations are deleted from the system.

See Also:

Oracle Grid Infrastructure on AIX, refer to *Oracle Grid Infrastructure Installation Guide for IBM AIX on POWER Systems (64-Bit)* for more information about installing

Note:

If a disk group is dismounted with the force option on AIX, any Oracle ADVM volumes device files that were on that disk group remain on the system. These files are removed when the disk group is remounted.

[Table 16-11](#) (page 16-42) lists the Oracle ACFS AIX commands with brief descriptions. The commands in [Table 16-11](#) (page 16-42) have been extended with additional options to support Oracle ACFS on AIX.

Table 16-11 Summary of Oracle ACFS commands for AIX

Command	Description
fsck (page 16-42)	Checks and repairs an Oracle ACFS file system on AIX.
mkfs (page 16-44)	Creates an Oracle ACFS file system on AIX.
mount (page 16-46)	Mounts an Oracle ACFS file system on AIX.
umount/umountall (page 16-48)	Dismounts an Oracle ACFS file system on AIX.

16.4.1 fsck

Purpose

Checks and repairs an Oracle ACFS file system on the AIX operating system.

Syntax and Description

```
fscck -V acfs [-n|-y] [-o options] volume_device
```

Table 16-12 (page 16-43) contains the options available with the AIX `fscck` command.

Table 16-12 Options for the AIX `fscck` command

Option	Description
<code>-V acfs</code>	Specifies an Oracle ADVM volume on AIX. <code>acfs</code> designates the Oracle ACFS type.
<code>-n</code>	Answers no to any prompts.
<code>-y</code>	Answers yes to any prompts.
<code>-o options</code>	Specifies that options follow (a, f, v). Options are preceded with the <code>-o</code> flag and entered as a comma-delimited string. For example: <code>-o a, v</code> <ul style="list-style-type: none"> • 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. Running in verbose mode can impact performance.
<code>volume_device</code>	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. Do not interrupt `fscck` during the repair operation.

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 incomplete 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 three such files are allocated. For example, a 2 GB file system being checked causes `fscck` to generate one to three 64K

working files in the `/usr/tmp` directory. These files are deleted after `fsck` has finished.

In the event that `fsck` 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 `fsck` is run in fix mode. For security reasons only the `root` user on Linux can 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 file](#) (page 16-147)".

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 16-11 Using the `fsck` command

```
# /usr/sbin/fsck -V acfs -y -o a /dev/asm/volume1-123
```

16.4.2 mkfs

Purpose

Creates an Oracle ACFS file system on the AIX operating system.

Syntax and Description

```
mkfs -V acfs -o h /dev/null
mkfs -V acfs [-v name ] [-s size][ -o options] volume_device
      [-c release_version]
```

`mkfs -V acfs -o h /dev/null` displays usage text and exits.

[Table 16-13](#) (page 16-45) contains the options available with the AIX `mkfs` command.

Table 16-13 Options for the AIX mkfs command

Option	Description
<code>-v acfs</code>	Specifies the type of file system on AIX. <code>acfs</code> designates the Oracle ACFS type.
<code>-v name</code>	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.
<code>-s size</code>	Specifies the size of the file system in 512-byte units or in units of K M G T P. Units specified are in K (kilobytes), M (megabytes), G (gigabytes), T (terabytes), or P (petabytes).
<code>-o options</code>	Specifies that options follow (f, h, v). Options are preceded with the <code>-o</code> flag and entered as a comma-delimited string. For example: <code>-o f,v</code> <ul style="list-style-type: none"> • <code>f</code> 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. Use this option with caution. • <code>h</code> Displays the usage text and exits. • <code>i=512 i=4096</code> Specifies a file system with metadata block size of 512 or 4096 bytes. • <code>v</code> Specifies verbose mode. The progress is displayed as the operation occurs.
<code>volume_device</code>	Specifies an Oracle ADVM device file that is to be formatted.
<code>-c release_version</code>	Creates an Oracle ACFS file system with Oracle ACFS release compatibility set to the value of <code>release_version</code> .

You can use `mkfs` to create the on disk structure needed for Oracle ACFS file system to be mounted. The `mkfs` command is the traditional Linux command used to build a file system. After `mkfs` runs 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 value specified with the `-c release_version` option must be greater than or equal to the value of `COMPATIBLE.ADMV` for the disk group, and must be less than or equal to the running Oracle Grid infrastructure release version. If `-c release_version` is not specified, the value of `COMPATIBLE.ADMV` is used. After the compatibility is set, it cannot be downgraded. The `-c release_version` option can be used in situations where it is not possible, or desirable, to update `COMPATIBLE.ADMV` and `COMPATIBLE.ASM` for the disk group, but you want to use an Oracle ACFS feature that requires a compatibility increase. After the Oracle ACFS compatibility is updated, it is no longer possible to mount the file system with an older Oracle Grid infrastructure release. To change the compatibility of an existing file system, refer to "[acfsutil compat set](#) (page 16-145)". For information about Oracle ASM disk group compatibility attributes, see "[Disk Group Compatibility](#) (page 4-50)".

The minimum file system size is 200 MB for the 512 byte metadata block size format and 512 MB for the 4 KB metadata block size format.

The Oracle ACFS driver must be loaded for `mkfs` to work.

4 KB Sector/Metadata

The following list summarizes the use of the `-o i` option and metadata block size based on the value of the `COMPATIBLE.ADVM` disk group attribute.

- If `COMPATIBLE.ADVM` is set to 12.2 or greater, then the metadata block size is 4096 bytes by default.
- If `COMPATIBLE.ADVM` is set to less than 12.2, then the block size is set to 512 bytes.
- If `COMPATIBLE.ADVM` is set to less than 12.2 and the logical sector size is not 512 bytes, then the command fails.

User data IO continues to support transfers as small as 512 bytes for normal user IO requests. When the Oracle ADVM volume of the file system has a logical disk sector size of 4 K, user Direct IO requests should be aligned on 4 K file offsets and be a multiple of 4 KB lengths for the best performance. Note that only 4 K metadata block size formats support ADVM volumes with 4 K logical disk sectors.

The `COMPATIBLE.ADVM` value must be set to 12.2 or higher to format an Oracle ADVM volume whose logical disk size is 4096 or to use the `-o i=4096` option to format an Oracle ADVM volume whose logical disk sector size is 512 bytes.

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 [Managing Oracle ADVM with ASMCMD](#) (page 15-1).

Next create an Oracle ACFS file system on the volume device file.

Example 16-12 Using the `mkfs` command

```
$ /usr/sbin/mkfs -V acfs /dev/asm/volume1-123
```

16.4.3 mount

Purpose

Mounts an Oracle ACFS file system on the AIX operating system.

Syntax and Description

```
mount -V acfs -o h /dev/null /xxx
mount -v acfs [-o options] volume_device mount_point
```

`mount -V acfs -o h /dev/null /xxx` displays usage text and exits.

Table 16-14 (page 16-47) contains the options available with the `mount` command.

Table 16-14 Options for the AIX mount command

Option	Description
<code>-v acfs</code>	Specifies the type of file system on AIX. <code>acfs</code> designates the Oracle ACFS type.
<code>-o options</code>	<p>Specifies that options follow. Options are preceded with the <code>-o</code> flag followed by a comma-delimited string of options. For example: <code>-o all,ro,v</code></p> <p>The following options are available:</p> <ul style="list-style-type: none"> • <code>all</code> Reads the Oracle ACFS managed file systems in the Oracle Clusterware namespace, mounting all file systems with the <code>AUTO_START=always</code>. The <code>-o all</code> option requires two placeholders to satisfy the command arguments: a dummy argument for the volume device and directory. When the <code>-o all</code> option is specified, other <code>-o</code> options are ignored. 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>h</code> Displays the usage help text and exits. • <code>nosuid</code> Specifies that the execution of <code>setuid</code> and <code>setgid</code> programs by way of this mount is not allowed. • <code>nodev</code> Specifies that you cannot open devices from this mount. • <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>ro</code> Mounts the file system in read-only mode. • <code>rw</code> Mounts the file system in read-write mode. • <code>v</code> Specifies verbose mode. The progress is displayed as the operation occurs.
<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>mount_point</code>	Specifies the directory where the file system is mounted. This directory must exist before you run the <code>mount</code> command.

`mount` attaches a file system to the Oracle ACFS hierarchy at the mount point that is the 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` runs successfully, the `MOUNTPATH` field in the `V$ASM_VOLUME` view displays the directory name on which the file system is now mounted.

An Oracle ACFS file system should only be mounted on one mount point. The same mount point name should be used on all cluster members.

`root` privilege is required to run `mount`.

Examples

The first example shows how to mount `volume1-123` on the mount point `/acfsmounts/acfs1`. The second example shows how to mount all the registered Oracle ACFS file systems. The dummy names (`none`) have been entered as placeholders for the device and directory as these arguments are required, but not used, when the `all` option is specified.

Example 16-13 Using the mount command

```
# /usr/sbin/mount -v acfs /dev/asm/volume1-123 /acfsmounts/acfs1
# /usr/sbin/mount -v acfs -o all none none
```

16.4.4 umount/umountall

Purpose

Dismounts an Oracle ACFS file system on the AIX operating system.

Syntax and Description

```
umount [mount_point | volume_device]
umountall -F acfs
```

`umountall -F acfs` dismounts all Oracle ACFS file systems.

[Table 16-15](#) (page 16-48) contains the options available with the `umount` command.

Table 16-15 Options for the AIX umount command

Option	Description
<code>mount_point</code>	Specifies the directory where the file system is mounted. Valid only with <code>umount</code> .
<code>volume_device</code>	Specifies the Oracle ADVM volume device name associated with the file system. Valid only with <code>umount</code> .
<code>-F acfs</code>	Specifies the type of file system on AIX. <code>acfs</code> designates the Oracle ACFS type. Valid only with <code>umountall</code> .

`umount` and `umountall` detach an Oracle ACFS from the file system hierarchy on the current node.

The `umount` and `umountall` commands do not return control back to the operating system prompt until the file system has been completely unmounted.

If a file system is busy, `umount` and `umountall` fail.

root privileges are required to run the `umount` and `umountall` commands.

Examples

The following examples show how to dismount an Oracle ACFS file system. The first example specifies the mount point of the file system to dismount. The second example specifies the volume device associated with the file system to dismount. The third example dismounts all Oracle ACFS file systems.

Example 16-14 Using the `umount` command

```
# /usr/sbin/umount /dev/asm/volume1-123

# /usr/sbin/umount /acfsmounts/acfs1

# /usr/sbin/umountall -F acfs
```

16.5 Oracle ACFS Command-Line Tools for Windows Environments

This topic provides a summary of the Oracle ACFS file system commands for Windows.

[Table 16-16](#) (page 16-49) lists the Oracle ACFS Windows file system commands with brief descriptions. These commands have been extended to support Oracle ACFS. All other Windows file system commands operate without change for Oracle ACFS.

Table 16-16 Summary of Oracle ACFS file system commands for Windows

Command	Description
acfschkdisk (page 16-50)	Checks and repairs an Oracle ACFS file system on Windows.
acfsdismount (page 16-52)	Dismounts an Oracle ACFS on Windows.
acfsformat (page 16-53)	Creates an Oracle ACFS file system on Windows.
acfsmountvol (page 16-54)	Mounts an Oracle ACFS file system on Windows.
acfsutil detach (page 16-56)	Prepares the Oracle ACFS driver for module unload on Windows.

[Table 16-17](#) (page 16-49) lists the non-Oracle ACFS Windows file system commands with brief descriptions.

Table 16-17 Summary of Non-Oracle ACFS file system commands for Windows

Command	Description
advmutil dismount (page 16-56)	Dismounts Oracle ADVM volume devices for file systems that are not Oracle ACFS on Windows.
advmutil list (page 16-58)	Lists Oracle ADVM volume devices for file systems that are not Oracle ACFS on Windows.

Table 16-17 (Cont.) Summary of Non-Oracle ACFS file system commands for Windows

Command	Description
advmutil mount (page 16-58)	Mounts Oracle ADVM volume devices for file systems that are not Oracle ACFS on Windows.

16.5.1 acfschkdsk

Purpose

Checks and repairs an Oracle ACFS file system.

Syntax and Description

```
acfschkdsk /h
acfschkdsk [/a|/f] [/v] [/n|/y] volume_device
```

[Table 16-18](#) (page 16-50) contains the options available with the `acfschkdsk` command.

Table 16-18 Options for the acfschkdsk command

Option	Description
/a	Specifies to automatically repair the file system.
/f	Forces the file system into mountable state without completing a file system check or repair.
/v	Specifies verbose mode. The progress is displayed as the operation occurs. Running in verbose mode can impact performance.
/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 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](#) (page 17-4)".

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 fail. 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 errors. The `/a` flag must be specified to instruct `acfschkdsk` to repair errors in the file system. Do not interrupt `acfschkdsk` during the repair operation.

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 incomplete 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 three 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 can 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 file](#) (page 16-147)".

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 16-15 Using the `acfschkdsk` command

```
C:\> acfschkdsk asm-volume1-311
```

16.5.2 acfsdismount

Purpose

Dismounts an Oracle ACFS file system and removes its mount point (name space attachment).

Syntax and Description

```
acfsdismount /h
acfsdismount [/v] {drive_letter | mount_point}
acfsdismount [/v] /all
```

[Table 16-19](#) (page 16-52) contains the options available with the `acfsdismount` command.

Table 16-19 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 ACFS file systems on this node.
<code>drive_letter</code>	Specifies the drive letter followed by a colon.
<code>mount_point</code>	Specifies the full path of the mount point including the drive letter.

`acfsdismount` removes the mount point (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.

`acfsdismount` can fail in some circumstances because there are open file references from various programs, such as a virus scanner, Process Explorer, or Oracle Enterprise Manager.

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 name. The third examples dismounts all mounted file systems. All the examples enable the verbose mode.

Example 16-16 Using the `acfsdismount` command

```
C:\> acfsdismount /v O:
C:\> acfsdismount /v O:\mnt\
C:\oracle>acfsdismount /v /all
acfsdismount: valid path name: C:\oracle\acfsmounts\acfs1\
acfsdismount: opening handle to volume C:\oracle\acfsmounts\acfs1
```

```
acfsdismount: locking volume C:\oracle\acfsmounts\acfs1
acfsdismount: dismounting volume C:\oracle\acfsmounts\acfs1
acfsdismount: removing volume mount point on C:\oracle\acfsmounts\acfs1\
acfsdismount: closing handle to volume C:\oracle\acfsmounts\acfs1 and exiting cleanly
```

16.5.3 acfsformat

Purpose

Creates an Oracle ACFS file system.

Syntax and Description

```
acfsformat /h
acfsformat [/v] [/i {512 | 4096}] [/f] [/n name] volume_device
           [/c release_version][blocks]
```

Table 16-20 (page 16-53) contains the options available with the `acfsformat` command.

Table 16-20 Options for the `acfsformat` command

Option	Description
<code>/v</code>	Specifies verbose mode. The progress is displayed as the operation occurs.
<code>/i {512 4096}</code>	Specifies a file system with metadata block size of 512 or 4096 bytes.
<code>/n name</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. Use this option with caution.
<code>/h</code>	Displays the usage help text.
<code>volume_device</code>	Specifies an Oracle ADVM device file to be formatted.
<code>/c release_version</code>	Creates an Oracle ACFS file system with Oracle ACFS release compatibility set to the value of <code>release_version</code> .
<code>blocks</code>	Specifies the number of blocks that the file system should consume on the named device. The quantity specified can be in units of <code>K</code> (kilobytes), <code>M</code> (megabytes), <code>G</code> (gigabytes), or <code>T</code> (terabytes). If a unit is not specified, the default is bytes. If this option is not specified, the entire device is consumed.

`acfsformat` creates the on disk structure needed for Oracle ACFS to be mounted. `acfsformat` performs a quick format on the target volume. After `acfsformat` runs successfully, the `USAGE` column in the `V$ASM_VOLUME` view displays ACFS.

The value specified with the `/c release_version` option must be greater than or equal to the value of `COMPATIBLE.ADVM` for the disk group, and must be less than or

equal to the running Oracle Grid infrastructure release version. If `/c release_version` is not specified, the value of `COMPATIBLE.ADVM` is used. After the compatibility is set, it cannot be downgraded. The `/c release_version` option can be used in situations where it is not possible, or desirable, to update `COMPATIBLE.ADVM` and `COMPATIBLE.ASM` for the disk group, but you want to use an Oracle ACFS feature that requires a compatibility increase. After the Oracle ACFS compatibility is updated, it is no longer possible to mount the file system with an older Oracle Grid infrastructure release. To change the compatibility of an existing file system, refer to “[acfsutil compat set](#) (page 16-145)”. For information about Oracle ASM disk group compatibility attributes, see “[Disk Group Compatibility](#) (page 4-50)”.

The minimum file system size is 200 MB for the 512 byte metadata block size format and 512 MB for the 4 KB metadata block size format.

The Oracle ACFS driver must be loaded for `acfsformat` to work.

Windows Administrator privileges are required to run this command.

4 KB Sector/Metadata

The following list summarizes the use of the `/i` option and metadata block size based on the value of the `COMPATIBLE.ADVM` disk group attribute.

- If `COMPATIBLE.ADVM` is set to 12.2 or greater, then the metadata block size is 4096 bytes by default.
- If `COMPATIBLE.ADVM` is set to less than 12.2, then the block size is set to 512 bytes.
- If `COMPATIBLE.ADVM` is set to less than 12.2 and the logical sector size is not 512 bytes, then the command fails.

User data IO continues to support transfers as small as 512 bytes for normal user IO requests. When the Oracle ADVM volume of the file system has a logical disk sector size of 4 K, user Direct IO requests should be aligned on 4 K file offsets and be a multiple of 4 KB lengths for the best performance. Note that only 4 K metadata block size formats support ADVM volumes with 4 K logical disk sectors.

The `COMPATIBLE.ADVM` value must be set to 12.2 or higher to format an Oracle ADVM volume whose logical disk size is 4096 or to use the `/i 4096` option to format an Oracle ADVM volume whose logical disk sector size is 512 bytes.

Examples

This example shows how to create an Oracle ACFS file system. The verbose option is specified in this example.

Example 16-17 Using the `acfsformat` command

```
C:\> acfsformat /v asm-volume1-311
```

16.5.4 `acfsmountvol`

Purpose

Mounts an Oracle ACFS file system.

Syntax and Description

```
acfsmountvol /h
acfsmountvol
acfsmountvol [/v] /all
acfsmountvol [/v] { drive_letter | dir } volume_device
```

Table 16-21 (page 16-55) contains the options available with the `acfsmountvol` command.

Table 16-21 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.
<code>drive_letter</code>	The Windows drive letter to link to an Oracle ADVM volume device.
<code>dir</code>	Specifies the full directory of the mount point including the drive letter.
<code>volume_device</code>	Specifies an Oracle ADVM volume device file that has been formatted by <code>acfsformat</code> .

`acfsmountvol` attaches an Oracle ACFS to the file system hierarchy at the specified path name or drive letter. `dir` must be an empty directory. Oracle ACFS mount points can be created on any empty directory and they can be hierarchical (nested).

After `acfsmountvol` runs successfully, the `MOUNTPATH` column in the `V $ASM_VOLUME` view displays the full path or drive letter on which the file system is now mounted. If no arguments are specified, the command lists all mounted Oracle ACFS file systems 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.

An Oracle ACFS file system should only be mounted on one mount point. The same mount point name should be used on all cluster members.

Windows Administrator privileges are required to mount an Oracle ACFS.

Examples

The first example shows how to mount `asm-volume1-311` on the mount point identified as the `F:` drive. The second example shows how to mount `asm-volume1-215` on the mount point specified by `c:\oracle\acfsmounts\acfs1\`. The first two examples enable verbose mode. The third example lists all the mounted file systems on the node.

Example 16-18 Using the `acfsmountvol` command

```
C:\> acfsmountvol /v F: asm-volume1-215
```

```
C:\> acfsmountvol /v c:\oracle\acfsmounts\acfs1 asm-volume1-311

C:\oracle>acfsmountvol
asm-volume1-311
C:\oracle\acfsmounts\acfs1
```

16.5.5 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 send 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 run 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 run 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 16-19 Using the acfsutil detach command

```
C:\> acfsdismount /all

C:\> acfsutil detach
```

16.5.6 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 | dir}
advmutil dismount [/unregister] [/all | /dg=disk_group]
```

[Table 16-22](#) (page 16-57) contains the options available with the Windows `advmutil dismount` command.

Table 16-22 Options for the `advmutil dismount` command

Option	Description
<code>drive_letter</code>	Specifies the drive letter to dismount.
<code>dir</code>	Specifies the directory of the mount point, including the drive letter.
<code>/all</code>	Specifies to dismount all file systems on this node.
<code>/unregister</code>	Optionally unregister an Oracle ADVM volume device and its drive letter from the system registry. 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 run 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.

The `advmutil dismount` command does not return control back to the operating system prompt until the file system has been completely unmounted.

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 16-20 Using the `advmutil dismount` command

```
C:\> advmutil dismount H:
C:\> advmutil dismount H: /unregister
C:\> advmutil dismount /all
```

```
C:\> advmutil dismount /dg=data
```

16.5.7 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 16-23](#) (page 16-58) contains the options available with the Windows `advmutil list` command.

Table 16-23 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 run 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 file](#) (page 16-147)".

Examples

The following example shows how to list all entries in the system registry.

Example 16-21 Using the `advmutil list` command

```
C:\> advmutil list /registry
```

```

Drive Letter      Asm Volume
=====
H:                asm-volume1-123
K:                asm-volume2-245
```

16.5.8 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 | dir} volume_device
advmutil mount /all
```

[Table 16-24](#) (page 16-59) contains the options available with the Windows `advmutil mount` command.

Table 16-24 Options for the *advmutil mount* command

Option	Description
<i>drive_letter</i>	Specifies the drive letter to mount.
<i>dir</i>	Specifies the directory of the mount point, including the drive letter.
<i>volume_device</i>	Indicates the Oracle ADVM volume device.
<i>/all</i>	Specifies to mount all non-Oracle ACFS file systems on this node.
<i>/register</i>	Optionally register 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 letters 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.

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 run 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](#) (page 16-160)".

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 ADVM volume device for use as a non-Oracle ACFS access which you could then use to format the volume for use with NTFS.
2. The second example shows how to create a drive letter for an Oracle ADVM volume device for use as a non-Oracle ACFS access which you could then use to format the volume 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 ADVM volume devices which were saved in the system registry which also allows a file system to be accessible.

Example 16-22 Using the *advmutil mount* command

```
C:\> advmutil mount H: asm-volume1-123

C:\> advmutil mount H: asm-volume1-123 /register

C:\> advmutil mount /all
```

16.6 Oracle ACFS Command-Line Tools for Tagging

This topic provides a summary of the Oracle ACFS tagging commands.

[Table 16-25](#) (page 16-60) lists the Oracle ACFS tagging commands with brief descriptions.

On Solaris, `acfsutil tag` commands can set tag names on symbolic link files, but backup and restore utilities do not save the tag names that are explicitly set on the symbolic link files. Also, symbolic link files lose explicitly set tag names if they have been moved, copied, tarred, or paxed.

Table 16-25 Summary of commands for Oracle ACFS tagging

Command	Description
acfsutil tag info (page 16-60)	Displays the tags for directories or files in Oracle ACFS file systems.
acfsutil tag set (page 16-61)	Adds a tag to directories or files in an Oracle ACFS file system.
acfsutil tag unset (page 16-63)	Removes a tag or all tag names from directories or files in an Oracle ACFS file system.

See Also:

- [Oracle ACFS Tagging](#) (page 11-39) for an overview of Oracle ACFS tagging
 - [About Using Oracle ACFS Command-Line Tools](#) (page 16-2) for information about running Oracle ACFS `acfsutil` commands
 - [Oracle ACFS Tagging Generic Application Programming Interface](#) (page 17-12) for information about Oracle ACFS tagging application programming interfaces (APIs)
-
-

16.6.1 acfsutil tag info

Purpose

Displays the tag names for tagged directories or file names in Oracle ACFS file systems.

Syntax and Description

```
acfsutil tag info -h
acfsutil tag info [-r] [-c -t tagname] path [path ...]
acfsutil tag info [-c -t tagname]
```

`acfsutil tag info -h` displays help text and exits.

[Table 16-26](#) (page 16-61) contains the options available with the `acfsutil tag info` command.

Table 16-26 Options for the `acfsutil tag info` command

Option	Description
<code>-t tagname</code>	Specifies the tag name string to display. Enclose the string in quotes if the string contains a space.
<code>path</code>	Specifies the path name to one or more files or directories.
<code>-r</code>	Specifies recursive action on a specified path.
<code>-c</code>	Specifies case-insensitive partial matching on the tag name.

The `acfsutil tag info` command can recursively apply the operation to all child files and subdirectories if the `-r` option is included with specified directory path names.

If no path names are specified, the `acfsutil tag info` command searches all Oracle ACFS mount points for tagged files.

Only the paths of tagged files in an Oracle ACFS file system are displayed. If the `-t tagname` option is included, only paths of files with a tag name exactly matching the specified `tagname` are displayed. If the `-c` option is included, then `tagname` can be a case-insensitive substring of a tag name for a successful match. For example, the `acfsutil tag info -c -t AG` command would display path names of files with tag names such as `tag1`, `ag`, or `AG`.

Any user may use this command to display tag info on a directory to which the user has read access.

Running `acfsutil tag info` on a file or directory that is protected by a security realm fails. If a user runs `acfsutil tag info` with the `-r` option on a directory protected by a security realm and the user does not have browse permissions on the directory, then no output is displayed. To resolve this issue, the user must be added to the security realm and granted `READDIR` permissions. For information about adding objects to an Oracle ACFS security realm, refer to "[acfsutil sec realm add](#)" (page 16-98).

Examples

The following example show the use of the `acfsutil tag info` command.

Example 16-23 Using the `acfsutil tag info` command

```
$ /sbin/acfsutil tag info -r /acfsmounts/acfs1/myrepfiles/
```

16.6.2 `acfsutil tag set`

Purpose

Adds the given tag to the specified files or directories in an Oracle ACFS file system

Syntax and Description

```
acfsutil tag set -h
acfsutil tag set [-v] [-r] tagname path [path ...]
```

`acfsutil tag set -h` displays help text and exits.

Table 16-27 (page 16-62) contains the options available with the `acfsutil tag set` command.

Table 16-27 Options for the `acfsutil tag set` command

Option	Description
<code>tagname</code>	Specifies a tag name string. Enclose the string in quotes if the string contains a space. The tag string can be composed of ASCII characters that include: a-b, A-Z, 0-9, the space character, the hyphen, and the underscore. The maximum length of a tag name is 32 characters. <code>all</code> is a reserved tag name and cannot be set on any file or directory.
<code>path</code>	Specifies the path string to one or more files or directories.
<code>-r</code>	Specifies recursive action on the specified path.
<code>-v</code>	Displays the progress of the operation.

The command can recursively apply the operation to all child files and subdirectories for the specified directory path names.

When adding a tag name to a file or directory, the existing tags on that file or directory remain unchanged. New files that are created after a directory is assigned a tag implicitly inherit tags from their parent directories. Existing files in the directory do not inherit the new tag; these files must be explicitly assigned the tag. Renaming a file or moving a file to a subdirectory within the same file system does not cause the file to inherit tag names from the new parent directory. Moving a file between file systems uses a copy operation and the newly created file does inherit the tag names of the parent directory.

There is no fixed maximum number of tag names for an Oracle ACFS file system or for each file and directory. However, the number of tag names is limited by the tag names that fit in up to 64 KB of special metadata storage for each file and directory. This metadata also contains information to manage the extended attributes. Longer tag names reduce the total of tag names that can be set on a file or directory. For example, if all tag names were 4 bytes long, then it is possible to hold approximately 1730 tag names, assuming no other extended attribute information is present.

Any user who has the privilege to modify the target file or directory may run this command.

If a user runs `acfsutil tag set` with the `-r` option on a directory protected by a security realm and does not have browse permissions on the directory, then no tags are applied. To resolve this issue, the user must be added to the security realm and granted `REaddir` permissions. For information about adding objects to an Oracle ACFS security realm, refer to "[acfsutil sec realm add](#) (page 16-98)".

Examples

The following example show the use of the `acfsutil tag set` command.

Example 16-24 Using the `acfsutil tag set` command

```
$ /sbin/acfsutil tag set repl_grpl -r /acfsmounts/acfs1/myrepfiles/*.dat
```

16.6.3 acfsutil tag unset

Purpose

Removes the given tag name or all tag names from the specified file or directory.

Syntax and Description

```
acfsutil tag unset -h
acfsutil tag unset [-v] [-r] {all | tagname} path [path ...]
```

acfsutil tag unset -h displays help text and exits.

[Table 16-28](#) (page 16-63) contains the options available with the `acfsutil tag unset` command.

Table 16-28 Options for the `acfsutil tag unset` command

Option	Description
<code>all</code>	Specifies to remove all tags in the path string.
<code>tagname</code>	Specifies a tag name string. Enclose the string in quotes if the string contains a space.
<code>path</code>	Specifies the path string to one or more files or directories.
<code>-r</code>	Specifies recursive action on the specified path.
<code>-v</code>	Displays the progress of the operation.

The command can recursively apply the operation to all child files and subdirectories for the specified directory path names.

When removing a tag name from a file or directory, other existing tags on that file or directory remain unchanged.

Any user who has privilege to modify the target file or directory may run this command.

If a user runs `acfsutil tag unset` with the `-r` option on a directory protected by a security realm and does not have browse permissions on the directory, then no tags are removed. To resolve this issue, the user must be added to the security realm and granted `READDIR` permissions. For information about adding objects to an Oracle ACFS security realm, refer to "[acfsutil sec realm add](#)" (page 16-98)".

Examples

The following example show the use of the `acfsutil tag unset` command.

Example 16-25 Using the `acfsutil tag unset` command

```
$ /sbin/acfsutil tag unset repl_grpl -r
  /acfsmounts/acfsl/myrepfiles/*.log
```

16.7 Oracle ACFS Command-Line Tools for Replication

This topic provides a summary of the commands for Oracle ACFS snapshot-based replication.

[Table 16-29](#) (page 16-64) lists the Oracle ACFS snapshot-based replication commands with brief descriptions. For an overview of Oracle ACFS replication, refer to [Oracle ACFS Replication](#) (page 11-36).

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Table 16-29 Summary of commands for Oracle ACFS replication

Command	Description
acfsutil repl bg (page 16-64)	Starts, stops, or displays information about the Oracle ACFS replication background processes.
acfsutil repl compare (page 16-65)	Verifies that files have been replicated on an Oracle ACFS file system.
acfsutil repl info (page 16-67)	Displays information about replication processing on an Oracle ACFS file system.
acfsutil repl init (page 16-69)	Initiates replication on files in an Oracle ACFS file system.
acfsutil repl pause (page 16-72)	Pauses replication on an Oracle ACFS file system.
acfsutil repl resume (page 16-73)	Resumes replication on an Oracle ACFS file system.
acfsutil repl sync (page 16-73)	Synchronizes primary and standby file systems.
acfsutil repl terminate (page 16-75)	Stops all replication activity on the Oracle ACFS file system.
acfsutil repl trace (page 16-76)	Specifies the trace level for replication information.
acfsutil repl update (page 16-76)	Updates replication information while replication is running.
acfsutil repl upgrade (page 16-78)	Upgrades an active, existing replication installation to use snapshot-based replication.

16.7.1 acfsutil repl bg

Purpose

Starts, stops, or displays information about the Oracle ACFS replication background processes.

Syntax and Description

```
acfsutil repl bg -h
acfsutil repl bg {start | stop | info} mount_point
```

`acfsutil repl bg -h` displays help text and exits.

[Table 16-30](#) (page 16-64) contains the options available with the `acfsutil repl bg` command.

Table 16-30 Options for the acfsutil repl bg command

Option	Description
<code>start</code>	Starts the replication background processing.
<code>stop</code>	Stops the replication background processing.

Table 16-30 (Cont.) Options for the `acfsutil repl bg` command

Option	Description
<code>info</code>	Displays information about the replication background process.
<code>mount_point</code>	Specifies the directory where the file system is mounted.

Oracle ACFS replication uses a background process on the primary cluster to transport file system changes to the standby file system. This process must be running for replication to function. When replication is started with the `acfsutil repl init` command, the process is started and registered with Oracle Clusterware, enabling an automatic restart after a reboot or system crash. This process is clusterwide.

To determine if the Oracle ACFS replication background process is running at the primary site, you can run the `acfsutil repl bg info` command on the primary cluster. For examples of the use of the command, refer to [Example 16-26](#) (page 16-65).

Note:

If `acfsutil repl bg info` shows that the replication background process is running, then the background process must be stopped before unmounting the Oracle ACFS file system. If a file system is participating in replication, it cannot be unmounted unless the Oracle ACFS replication background process is first stopped. Any attempt to unmount the file system without first stopping this process fails. System shutdown automatically runs the `acfsutil repl bg` command before unmounting ACFS file systems.

You must have system administrator privileges to run the `acfsutil repl bg start` or `acfsutil repl bg stop` commands. You must have system administrator or Oracle ASM administrator privileges to run the `acfsutil repl bg info` command.

Examples

[Example 16-26](#) (page 16-65) shows the use of the `acfsutil repl bg` command.

Example 16-26 Using the `acfsutil repl bg` command

```
$ /sbin/acfsutil repl bg start /acfsmounts/acfs1

$ /sbin/acfsutil repl bg info /acfsmounts/acfs1
ACFS clusterwide replication background process is running.

$ /sbin/acfsutil repl bg stop /acfsmounts/acfs1

$ /sbin/acfsutil repl bg info /acfsmounts/acfs1
ACFS clusterwide replication background process is stopped.
```

16.7.2 acfsutil repl compare**Purpose**

Verifies that files have been replicated on an Oracle ACFS file system.

Syntax and Description

```
acfsutil repl compare -h
acfsutil repl compare [-v] [ -a [-s] | -t { all | tagname, ... } [-s] ]
primary_mount_point standby_mount_point
```

`acfsutil repl compare -h` displays help text and exits.

[Table 16-31](#) (page 16-66) contains the options available with the `acfsutil repl compare` command.

Table 16-31 Options for the `acfsutil repl compare` command

Option	Description
<code>-a</code>	Compares all files and their extended attribute names.
<code>-t { all tagname, ... }</code>	Compares files with all or some tag names specified in the <code>acfsutil repl init</code> operation. The <code>all</code> option specifies all tag names supplied in the <code>acfsutil repl init</code> operation. A comma-delimited list of tag names must only use tag names specified during the <code>acfsutil repl init</code> operation.
<code>-s</code>	Skips comparisons on symlinks for extended attributes and tags when using the <code>-a</code> or <code>-t</code> option. The symlinks are still checked for everything else, such as permissions and ownership.
<code>primary_mount_point</code>	Specifies the directory where the primary file system is mounted.
<code>standby_mount_point</code>	Specifies the directory where the standby file system is mounted.

The `acfsutil repl compare` command verifies all or part of the primary mount point has been replicated to the standby mount point. The standby mount point must be mounted locally for comparison. Best results are obtained when the primary file system is changing as little as possible. The command always checks all files on the primary file system against those on the standby file system.

The `-a` option is typically used when no tag names were specified during the `acfsutil repl init` operation. The `-a` option compares all files on the primary file system against those on the standby file system. The `-a` option also enables tests for extra files on the standby file system that do not exist on the primary. When testing for extra files, the command also attempts to check extended attributes. However, if NFS is used to mount the standby file system locally, the standby cannot be checked for matching extended attributes due to limitations in the NFS protocol.

The `-t` option is used when tags were specified during the `acfsutil repl init` operation. The `-t` operation locates all file names on the primary file system with the specified tag names and compares them to the corresponding files on the standby. The `-t` option takes as arguments either `all` option or a comma-delimited list of tag names. If the `-t` option is used with the `all` argument, then all tag names supplied during the `acfsutil repl init` operation are selected. Otherwise, any specified tag names must exactly match a tag name supplied during the `acfsutil repl init` operation. The `-t` option also tests for extra files on the standby file system that do not have an associated tag name that was specified during the `acfsutil repl init` operation. If NFS is used to mount the standby file system locally, the standby cannot be checked for matching tag names and extended attributes due to limitations in the NFS protocol.

You can use the `acfsutil repl info -c` option to determine which tags were specified during the `acfsutil repl init` operation.

The `-a` and `-t` options cannot both be specified at that same time. When `-a` or `-t` is provided, the `-s` option can also be specified to skip extended attribute comparisons for symlinks. If neither the `-a` nor `-t` option is provided, a primary to standby file comparison is done without testing tag names nor extended attributes.

The `-v` option may be specified to output the name of each file after the file is compared.

You must have system administrator or Oracle ASM administrator privileges to run this command.

Examples

[Example 16-27](#) (page 16-67) shows the use of the `acfsutil repl compare` command.

Example 16-27 Using the `acfsutil repl compare` command

```
$ /sbin/acfsutil repl compare /acfsmounts/acfs1 /nfs_mounted_standby
```

16.7.3 acfsutil repl info

Purpose

Displays information about replication processing on an Oracle ACFS file system.

Syntax and Description

```
acfsutil repl info -h
```

```
acfsutil repl info -c [-v] mount_point
```

```
acfsutil repl info -c -u user [-o sshStrictKey=yvalue] standby_host
[standby_host]...]
    primary_mount_point
```

```
acfsutil repl info -s [-v][-l] [-n number{m|h|d|w|y}]
    {-f eventlog | mount_point}
```

```
acfsutil repl info [-a|-e|-t] [-v][-l]
    {-f eventlog | mount_point }
```

`acfsutil repl info -h` displays help text and exits.

[Table 16-32](#) (page 16-67) contains the options available with the `acfsutil repl info` command.

Table 16-32 Options for the `acfsutil repl info` command

Option	Description
<code>-a</code>	Displays only apply records from the replication event log. An apply record contains the date and time that the set of file system changes were captured on the primary file system and the date and time that they were applied on the standby file system.

Table 16-32 (Cont.) Options for the `acfsutil repl info` command

Option	Description
<code>-c</code>	Displays configuration information.
<code>-e</code>	Displays only error records from the replication event log.
<code>-f event_log</code>	Specifies the absolute path to the replication event log to be used as the source of the information. The default location is: <code>mount_point/.ACFS/repl/logs/ReplicationEventsLog</code>
<code>-h</code>	Displays help text.
<code>-l</code>	Displays last event only.
<code>-n number{m h d w y}</code>	Specifies the interval to display information. <i>number</i> specifies the number of minutes (m), hours (h), days (d), weeks (w), or years (y). For example: <code>-n 5d</code> The statistics are summarized in the units of time specified.
<code>-u user</code>	Validates the host and user keys currently configured between the current host and each standby host specified with a <i>standby</i> argument, using <i>user</i> as the user for which access is validated.
<code>-o sshStrictKey=yvalue</code>	Specifies whether <code>ssh</code> should use strict host-key checking. A value starting with <code>y</code> enables this checking, which is the default setting. A value starting with <code>n</code> disables the checking.
<code>-s</code>	Displays statistical information. This option is valid only when run on the primary file system. It is important that the primary and standby nodes are running a network time service to ensure that the statistics are meaningful.
<code>-t</code>	Displays only transport records from the replication event log.
<code>-v</code>	Displays verbose output.
<i>standby_host</i>	Specifies a hostname or IP address for one of the standby hosts to be processed when the <code>-u</code> option is specified.
<i>primary_mount_point</i>	Specifies the directory where the primary file system is mounted.
<i>mount_point</i>	Specifies the directory where the file system is mounted (primary or standby).

To display information about replication configuration, use `acfsutil repl info` with the `-c` option. The configuration information includes a `Applied primary as of` timestamp that identifies the point in time on the primary file system that the standby file system represents. You can use this timestamp to set back a database to this point in time if it is necessary to keep the data in Oracle ACFS file system synchronized with the database data replicated with Oracle Data Guard.

When specifying the `-u` option, the command should be run as `root` (or `local SYSTEM`) on each host of the primary cluster to confirm that `ssh`, as used by snapshot-based replication, can connect to all standby hosts and log in as *user* on each host. The *user* specified should be the same as the *repluser* specified in the `acfsutil repl init`

command. Each *standby_host* string specifies a hostname or host IP address for a standby node. The specified *primary_mount_point* should be the active Oracle ACFS mount point in the primary cluster that is to be used for replication. If strict host-key checking was or is to be disabled for this use of `acfsutil repl init primary` using `-o sshStrictKey=no`, then the option should be specified here also to skip host-key validation. `acfsutil repl info -c -u` can be run at any time, whether or not replication is currently active on the *primary_mount_point*.

To display information about replication statistics, use `acfsutil repl info` with the `-s` option.

To display information about all replication events, use `acfsutil repl info` with no options. To display information about a specific kind of event, use `acfsutil repl info` with the `-a`, `-e`, or `-t` option.

You must have system administrator or Oracle ASM administrator privileges to run this command.

Examples

[Example 16-28](#) (page 16-69) shows the use of the `acfsutil repl info` command.

Example 16-28 Using the `acfsutil repl info` command

The validation command has the following format:

```
# /sbin/acfsutil repl info -c -u repluser standby-addr1 [standby-addr2 ...]
my_primary_mountpoint
```

The command confirms that user `repluser` can use `ssh` to connect to each *standby-addr* given in the same way as replication does when initializing. The *my_primary_mountpoint* given should be the active Oracle ACFS mount point that is to be used for replication.

To validate the key setup for the *standby1* and *standby2* cluster presented in the previous command, you can use the following command on each node of the primary cluster:

```
# /sbin/acfsutil repl info -c -u repluser standby1 standby2 my_primary_mountpoint
```

The same command would be used if you were planning to use the VIP *standby12_vip* to connect to the cluster.

If you plan to disable strict host key checking, you can skip this checking by adding the `-o sshStrictKey=no` option to the command line.

Additional commands for displaying statistics and apply records are the following:

```
$ /sbin/acfsutil repl info -s -n 5d /acfsmounts/acfs1
```

```
$ /sbin/acfsutil repl info -a -v /acfsmounts/acfs1
```

16.7.4 `acfsutil repl init`

Purpose

Initiates replication on all the files in an Oracle ACFS file system or only those files with a specified list of tags.

Syntax and Description

```
acfsutil repl init -h
```

```
acfsutil repl init standby -u repluser [-d trace_level] standby_mount_point
```

```
acfsutil repl init primary -s repluser@netname { -C | -i interval } [tagname...]
[-m standby_mount_point] [-d trace_level] [-z {on | off}]
[-o sshCmdPath=pathname] [-o sshStrictKey=yvalue]
[-o sshCipher=ciphername]
primary_mount_point
```

acfsutil repl init -h displays help text and exits.

[Table 16-33](#) (page 16-70) contains the options available with the `acfsutil repl init` command.

Table 16-33 Options for the `acfsutil repl init` command

Option	Description
primary	Specifies replication of the primary file system.
standby	Specifies replication of the standby file system.
-s repluser@netname	Specifies a standard <code>ssh</code> connection string for the primary site to contact the standby site. The <i>repluser</i> in the connection string is the user that <code>ssh</code> logs in as on the standby node. Choose a minimally-privileged user identity. The user should have Oracle ASM administrator privileges and should belong to the Oracle DBA group. The <i>netname</i> in the connection string specifies a network endpoint, such as a host name, VIP name, or IP address.
-C	Specifies replication in constant mode. Snapshots are continuously generated and replicated to the standby site. As soon as the replication of each snapshot completes, the generation of a new snapshot is started.
-i interval	Specifies replication in interval (scheduled) mode. A new snapshot is taken and replicated with the frequency specified, if possible. A suffix must be given to specify the units in which interval is measured. The suffix must be either <i>s</i> (seconds), <i>m</i> (minutes), <i>h</i> (hours), <i>d</i> (days), or <i>w</i> (weeks). For example, 30m is thirty minutes and 2h is two hours.
tagname	Specifies one or more tagnames to be used in restricting replication. When tagnames are specified, only files that are tagged with at least one of the those tagnames are replicated. For information about Oracle ACFS tags, refer Oracle ACFS Tagging (page 11-39).
-m standby_mount_point	Specifies the directory where the standby file system is mounted if it is different than the primary file system.
-d trace_level	Specifies the trace level setting [0..6].
-z on off	Turns on or off compression of the replication data stream sent from primary to standby.

Table 16-33 (Cont.) Options for the `acfsutil repl init` command

Option	Description
<code>-o</code> <code>sshCmdPath=pathname</code>	Specifies the pathname to the <code>ssh</code> command.
<code>-o</code> <code>sshStrictKey=y/nvalue</code>	Specifies whether <code>ssh</code> should use strict host-key checking. A value starting with <code>y</code> enables this checking, which is the default setting. A value starting with <code>n</code> disables the checking.
<code>-o</code> <code>sshCipher=ciphername</code>	Specifies the cipher that is passed to <code>ssh</code> to encrypt its sessions.
<code>-u</code>	Specifies the <code>repluser</code> .
<code>standby_mount_point</code>	Specifies the directory where the standby file system is mounted.
<code>primary_mount_point</code>	Specifies the directory where the primary file system is mounted.

The `acfsutil repl init primary` and `standby` commands initiate replication between a specified primary site and a specified standby site. The `acfsutil repl init standby` command must be run first. Before running the command, ensure that the file system specified is mounted on only one node of the standby cluster. When the command completes successfully, mount the specified file system on all nodes of the standby cluster. Next, run the `acfsutil repl init primary` command.

When the `acfsutil repl init primary` command completes successfully, processing begins to replicate initial copies of all specified files to the standby file system. In addition, any changes to these files and any new files subsequently created in the file system on the primary file system are captured and transported to the standby file system. However, successful completion of the command does not indicate that the specified files have been copied in full to the standby file system. To monitor the progress of the initial replication of all specified files to the standby file system, users can run the `acfsutil repl info -c` command.

The initial replication operation is based on a snapshot of the primary file system, which is transferred to the standby file system in its entirety. Each subsequent replication operation uses a later snapshot of the primary, and only the differences between that snapshot and the last previous snapshot are transferred to the standby.

All directories in the primary file system are always replicated to the standby, even if tags are specified.

Symbolic link files are replicated as is. If the symbolic link resolves to an absolute path name, and that path name does not exist on the standby file system, then referencing the symbolic link results in errors.

On platforms that support named sockets, character device files or block device files, these file types are not replicated.

A file system being replicated can be mounted on only one path for the primary or standby file system.

A replicated Oracle ACFS file system must be mounted on the same directory within a cluster. The file system does not have to be mounted on all of the nodes in a cluster, but if it is mounted, it must be mounted on the directory specified with the `acfsutil`

`repl init` command. In addition, no other file system can be mounted on that directory on other nodes.

A single Oracle ACFS file system cannot be configured both as a primary and a standby file system. If replication is configured on the same host or within the same cluster, then the primary and standby file systems cannot use the same named mount point.

If either site hosting the primary or standby file system is running on Windows, then both sites must be running on Windows. Replication is allowed between sites running any combination of the Linux, Solaris or AIX operating systems.

You must have system administrator privileges to run this command.

Examples

[Example 16-29](#) (page 16-72) shows the use of the `acfsutil repl init` to initiate snapshot-based replication from the primary to standby site.

Example 16-29 Using the `acfsutil repl init` command

```
# /sbin/acfsutil repl init standby -u my_repluser /my_standby/repl_data

# /sbin/acfsutil repl init primary -s my_repluser@my_host -i 2h
                                     -m /my_standby/repl_data /my_primary/repl_data
```

16.7.5 `acfsutil repl pause`

Purpose

Pauses replication on an Oracle ACFS file system.

Syntax and Description

```
acfsutil repl pause -h
acfsutil repl pause mount_point
```

`acfsutil repl pause -h` displays help text and exits.

[Table 16-34](#) (page 16-72) contains the options available with the `acfsutil repl pause` command.

Table 16-34 Options for the `acfsutil repl pause` command

Option	Description
<code>mount_point</code>	Specifies the directory where the file system is mounted.

The `acfsutil repl pause` command is issued at the primary replication site only. To continue replication operations, `acfsutil repl resume` should be run at a later time on the same primary site.

`acfsutil repl pause` stops the operation of the replication daemon temporarily. Subsequent running of `acfsutil repl resume` resumes the operation of the daemon.

If `acfsutil repl pause` is run while a replication operation is active, then the active replication operation completes regardless of the pause action. The pause operation is not effective in this case, and the command does not return, until the current data

stream has been applied on the standby. After that, operations are not performed for this replication until the `acfsutil repl resume` is issued.

You must have system administrator or Oracle ASM administrator privileges to run this command.

Examples

[Example 16-30](#) (page 16-73) shows the use of the `acfsutil repl pause` command.

Example 16-30 Using the `acfsutil repl pause` command

```
$ /sbin/acfsutil repl pause /acfsmounts/acfs1
```

16.7.6 `acfsutil repl resume`

Purpose

Resumes replication on an Oracle ACFS file system where replication has been paused.

Syntax and Description

```
acfsutil repl resume -h
acfsutil repl resume mount_point
```

`acfsutil repl resume -h` displays help text and exits.

[Table 16-35](#) (page 16-73) contains the options available with the `acfsutil repl resume` command.

Table 16-35 Options for the `acfsutil repl resume` command

Option	Description
<code>mount_point</code>	Specifies the directory where the file system is mounted.

The `acfsutil repl resume` command resumes the replication process after the `acfsutil repl pause` command has been run. This command should be run on the same primary site soon after replication has been paused with the `acfsutil repl pause` command.

You must have system administrator or Oracle ASM administrator privileges to run this command.

Examples

[Example 16-31](#) (page 16-73) shows the use of the `acfsutil repl resume` command.

Example 16-31 Using the `acfsutil repl resume` command

```
$ /sbin/acfsutil repl resume /acfsmounts/acfs1
```

16.7.7 `acfsutil repl sync`

Purpose

Synchronizes primary and standby file systems.

Syntax and Description

```
acfsutil repl sync -h
acfsutil repl sync [apply] primary_mount_point
```

`acfsutil repl sync -h` displays help text and exits.

[Table 16-36](#) (page 16-74) contains the options available with the `acfsutil repl sync` command.

Table 16-36 Options for the `acfsutil repl sync` command

Option	Description
<code>apply</code>	Specifies that all changes be applied to the standby file system. Currently has no effect on the operation of the command.
<code>primary_mount_point</code>	Specifies the directory where the primary file system is mounted.

The `acfsutil repl sync` command can be used to synchronize the state of the primary site and the standby site. This command can only be run on the primary site. Users should first quiesce their applications and issue a `sync(1)` call to ensure that the synchronized state is known and meaningful to the user. Running `acfsutil repl sync` then ensures all outstanding replication data is shipped from the primary site to the standby site. Specifically, the command enables any in-progress replication operation to complete, and then performs replication one final time to ensure that all changes on the primary have been replicated.

The command returns success when all of the changes have been successfully applied to the file system on the standby site. At this point, unless an unmount of the last node on the primary site cluster is about to occur, applications may be restarted.

Successfully unmounting a replicated file system on the primary site does not guarantee that all file system changes made prior to the unmount have been successfully sent to the standby site or applied to the standby file system. If the primary file system is unmounted on one primary site node but remains mounted on one or more primary site nodes, changes to the file system made before the unmount continue to be transported to the standby site from the other nodes after the unmount. However, if you are in the process of unmounting the primary file system on the last primary site node where it is mounted, and if you want to know that all file system changes on the primary file system up to that point have been successfully applied to the standby file system, then you should perform the following operations to ensure that the standby file system is up to date.

- Quiesce applications on the primary site which modify the file system.
- Run `acfsutil repl sync`.

Next, run `acfsutil repl bg stop` to stop the replication background process, which always must be run before unmounting a replicated file system. After the background process has been stopped, then unmount the file system.

You must have system administrator or Oracle ASM administrator privileges to run this command.

Examples

[Example 16-32](#) (page 16-75) shows the use of the `acfsutil repl sync` command.

Example 16-32 Using the `acfsutil repl sync` command

```
$ /sbin/acfsutil repl sync /acfsmounts/acfs1
```

16.7.8 acfsutil repl terminate**Purpose**

Stops all replication activity on the Oracle ACFS file system at the site where it is run.

Syntax and Description

```
acfsutil repl terminate -h
acfsutil repl terminate primary [-d trace_level] mount_point
acfsutil repl terminate standby [immediate] [-k] [-d trace_level] mount_point
```

`acfsutil repl terminate -h` displays help text and exits.

Table 16-37 (page 16-75) contains the options available with the `acfsutil repl terminate` command.

Table 16-37 Options for the `acfsutil repl terminate` command

Option	Description
<code>primary</code>	Stops replication on the primary file system.
<code>standby</code>	Stops replication on the standby file system.
<code>immediate</code>	Stops replication processing immediately on the standby file system.
<code>-k</code>	Specifies to keep the backup snapshot maintained for the standby. Without this option, the snapshot is deleted.
<code>-d trace_level</code>	Specifies the trace level setting [0..6].
<code>mount_point</code>	Specifies the directory where the file system is mounted.

The `acfsutil repl terminate` command stops all replication processing. When terminating replication, you should terminate replication for the primary file system first, and then the standby file system. If you want to ensure all changes are sent to the standby file system before terminating the primary site, ensure that all applications are quiesced and run the `acfsutil repl sync` command. The terminate command must be run on both sites if both sites hosting the file systems are available.

After `acfsutil repl terminate standby` has been run, you can use the standby file system in read-write mode. If you want to restart replication after running `acfsutil repl terminate`, then you must restart replication from the beginning with the `acfsutil repl init` command.

The `repl terminate standby` command waits until replication data at the standby file system has been applied. If you want to terminate replication immediately without applying all the replication data, use the `immediate` option. However, this option can leave some files at the standby file system in an indeterminate state.

To momentarily stop Oracle ACFS replication, you can use the `acfsutil repl pause` command, followed soon after by the `acfsutil repl resume` command.

You must have system administrator privileges to run this command.

Examples

[Example 16-33](#) (page 16-76) shows the use of the `acfsutil repl terminate` command.

Example 16-33 Using the `acfsutil repl terminate` command

```
$ /sbin/acfsutil repl terminate /acfsmounts/acfs1
```

16.7.9 acfsutil repl trace**Purpose**

Sets the replication trace level for gathering trace information on an Oracle ACFS file system.

Syntax and Description

```
acfsutil repl trace -h
acfsutil repl trace level mount_point
```

`acfsutil repl trace -h` displays help text and exits.

[Table 16-39](#) (page 16-77) contains the options available with the `acfsutil repl trace` command.

Table 16-38 Options for the `acfsutil repl trace` command

Option	Description
<i>level</i>	Specifies the trace level setting [0..6]. The default level is 2.
<i>mount_point</i>	Specifies the directory where the file system is mounted.

Increasing the trace level can have a performance impact and should be done at the recommendation of Oracle support services.

Trace files for Oracle ACFS replication are stored in the following location:

```
GRID_HOME/diag/crs/hostname/crs/trace
```

You must have system administrator or Oracle ASM administrator privileges to run this command.

Examples

[Example 16-34](#) (page 16-76) shows the use of the `acfsutil repl trace` command.

Example 16-34 Using the `acfsutil repl trace` command

```
$ /sbin/acfsutil repl trace 5 /acfsmounts/acfs1
```

16.7.10 acfsutil repl update**Purpose**

Updates replication information after replication has started and is running on an Oracle ACFS file system.

Syntax and Description

```
acfsutil repl update -h

acfsutil repl update [-s repluser@netname] [-C | -i interval]
                    [-d trace_level] [-z {on | off}]
                    [-o sshCmdPath=pathname] [-o sshStrictKey=yvalue] [-o
sshCipher=ciphername]
                    primary_mount_point

acfsutil repl update [-u repluser] [-d trace_level] standby_mount_point

acfsutil repl update -h displays help text and exits.
```

[Table 16-39](#) (page 16-77) contains the options available with the `acfsutil repl update` command.

Table 16-39 Options for the `acfsutil repl update` command

Option	Description
<code>-s repluser@netname</code>	Specifies a standard <code>ssh</code> connection string for the primary site to contact the standby site. The <i>repluser</i> in the connection string is the user that <code>ssh</code> logs in as on the standby node. Choose a minimally-privileged user identity. The user should have Oracle ASM administrator privileges and should belong to the Oracle DBA group. The <i>netname</i> in the connection string specifies a network endpoint, such as a host name, VIP name, or IP address.
<code>-C</code>	Specifies replication in constant mode. Snapshots are continuously generated and replicated to the standby site. As soon as the replication of each snapshot completes, the generation of a new snapshot is started.
<code>-i interval</code>	Specifies replication in interval (scheduled) mode. A new snapshot is taken and replicated with the frequency specified, if possible. A suffix must be given to specify the units in which interval is measured. The suffix must be either <i>s</i> (seconds), <i>m</i> (minutes), <i>h</i> (hours), <i>d</i> (days), or <i>w</i> (weeks). For example, 30m is thirty minutes and 2h is two hours.
<code>-d trace_level</code>	Specifies the trace level setting [0..6].
<code>-z on off</code>	Turns on or off compression of the replication data stream sent from primary to standby.
<code>-o sshCmdPath=pathname</code>	Specifies the pathname to the <code>ssh</code> command.
<code>-o sshStrictKey=yvalue</code>	Specifies whether <code>ssh</code> should use strict host-key checking. A value starting with <i>y</i> enables this checking, which is the default setting. A value starting with <i>n</i> disables the checking.
<code>-o sshCipher=ciphername</code>	Specifies the cipher that is passed to <code>ssh</code> to encrypt its sessions.
<code>-u</code>	Specifies the <i>repluser</i> .
<code>primary_mount_point</code>	Specifies the directory where the primary file system is mounted.

The `acfsutil repl update` command updates replication information and configuration after replication has been started. For example, a system administrator can change the network interface that is currently being used by replication to connect to the remote site. Issuing the `acfsutil repl update` command to update the configuration in this situation enables replication to continue running uninterrupted.

If `acfsutil repl update` is run on the primary site, the command can be used to alter the use `rname` and network interface name (hostname or VIP name) that is used to connect to the standby site. The command can also be used to change the replication interval or tracing level, change the cipher used by or the pathname used for `ssh`, and turn compression on or off.

If `acfsutil repl update` is run on the standby site, the command can be used to alter either the tracing level or the user name used by the primary site to connect to the standby site.

At least one of the options must be specified when issuing this command; the command cannot be run with only the required mount point.

The `acfsutil repl update` command returns success when the updated information has been accepted on the local site. To change the user name, the command must be run on both the primary and the standby sites. If the replication interval is changed with the `-C` or `-i` option, a replication operation occurs at the point when `acfsutil repl update` is run. The next replication operation occurs based on the newly-specified interval.

You must have system administrator privileges to run this command.

Examples

[Example 16-35](#) (page 16-78) shows the use of the `acfsutil repl update` command.

Example 16-35 Using the `acfsutil repl update` command

```
$ /sbin/acfsutil repl update -i lh /my_primary/repl_data
```

16.7.11 `acfsutil repl upgrade`

Purpose

Upgrades an existing replication installation to snapshot-based replication.

Syntax and Description

```
acfsutil repl upgrade -h

acfsutil repl upgrade prepare -s {repluser@netname | -u} primary_mount_point

acfsutil repl upgrade standby -u repluser standby_mount_point

acfsutil repl upgrade primary {-C | -i interval}
                               [-d trace_level] [-z {on | off}]
                               [-o sshCmdPath=pathname] [-o sshStrictKey=yynvalue]
                               [-o sshCipher=ciphername]
                               primary_mount_point
```

`acfsutil repl upgrade -h` displays help text and exits.

[Table 16-39](#) (page 16-77) contains the options available with the `acfsutil repl upgrade` command.

Table 16-40 Options for the `acfsutil repl upgrade` command

Option	Description
<code>prepare</code>	Prepares an existing replication installation for upgrade to snapshot-based replication.
<code>-srepluser@netname</code>	Specifies a standard <code>ssh</code> connection string for the primary site to contact the standby site. The <code>repluser</code> in the connection string is the user that <code>ssh</code> logs in as on the standby node. Choose a minimally-privileged user identity. The user should have Oracle ASM administrator privileges and should belong to the Oracle DBA group. The <code>netname</code> in the connection string specifies a network endpoint, such as a host name, VIP name, or IP address.
<code>-u (prepare option)</code>	Undoes a completed <code>acfsutil repl upgrade prepare</code> operation. This action restores the <code>primary_mount_point</code> to its state before the previously completed <code>acfsutil repl upgrade prepare</code> operation.
<code>standby</code>	Specifies replication of the standby file system.
<code>-u (standby option)</code>	Specifies the <code>repluser</code> .
<code>standby_mount_point</code>	Specifies the directory where the standby file system is mounted.
<code>primary</code>	Specifies replication of the primary file system.
<code>-C</code>	Specifies replication in constant mode. Snapshots are continuously generated and replicated to the standby site. As soon as the replication of each snapshot completes, the generation of a new snapshot is started.
<code>-i interval</code>	Specifies replication in interval (scheduled) mode. A new snapshot is taken and replicated with the frequency specified, if possible. A suffix must be given to specify the units in which interval is measured. The suffix must be either <code>s</code> (seconds), <code>m</code> (minutes), <code>h</code> (hours), <code>d</code> (days), or <code>w</code> (weeks). For example, <code>30m</code> is thirty minutes and <code>2h</code> is two hours.
<code>-d trace_level</code>	Specifies the trace level setting (0..6).
<code>-z on off</code>	Turns on or off compression of the replication data stream sent from primary to standby.
<code>-o sshCmdPath=pathname</code>	Specifies the pathname to the <code>ssh</code> command.
<code>-o sshStrictKey=y nvalue</code>	Specifies whether <code>ssh</code> should use strict host-key checking. A value starting with <code>y</code> enables this checking, which is the default setting. A value starting with <code>n</code> disables the checking.
<code>-o sshCipher=ciphername</code>	Specifies the cipher that is passed to <code>ssh</code> to encrypt its sessions.
<code>primary_mount_point</code>	Specifies the directory where the primary file system is mounted.

The `acfsutil repl upgrade` command provides support for upgrading an active, existing replication relationship to use snapshot-based replication. The intent of the upgrade process is to provide an active snapshot-based replication relationship between the current primary and standby file systems without terminating or reinitiating replication and without losing any data in the replicating process.

Before you start the upgrade procedure, you must ensure that the primary file system is mounted on only one node of its cluster. You should quiesce any application activity targeting the primary or standby file system to ensure an optimal upgrade; however, the upgrade proceeds if this recommendation is ignored.

Note:

After the upgrade has been started, the process must be completed. The `acfsutil repl upgrade prepare` operation can be undone using the `-u` option. However, after the `acfsutil repl upgrade standby` operation has been started, the entire upgrade process must be completed. From that point forward, there is no provision for stopping or undoing the upgrade.

To begin the upgrade process, first run the `acfsutil repl upgrade prepare` command on the primary cluster. This command specifies the user and host or interface name, as well as the primary mount point, that is used for snapshot-based replication. The user and host names are given with the `-s` option, the same as with the `acfsutil repl init primary` command option for snapshot-based replication.

Next, upgrade the standby site. On the standby, run the `acfsutil repl upgrade standby` command. This command specifies the `repluser`, as well as the standby mount point, that is used for snapshot-based replication. The user name is specified with the `-u` option, the same as with the `acfsutil repl init standby` command option for snapshot-based replication.

After this command has been issued, the upgrade must be completed by issuing the `acfsutil repl upgrade primary` command.

After the `acfsutil repl upgrade standby` command has been run on the standby site, the `acfsutil repl upgrade primary` command must be run on the primary site to complete the upgrade to snapshot-based replication. This command automatically terminates the original replication process and initiates snapshot-based replication. `acfsutil repl upgrade primary` accepts any of the command-line options accepted by the `acfsutil repl init primary` command for snapshot-based replication, except for the `-m` option, `-s` option, and tag names. The necessary information for the `-m` and `-s` options has been obtained from other sources. Any tag names that were specified for the original replication relationship are automatically transferred to the new relationship. For information about Oracle ACFS tags, refer to [Oracle ACFS Tagging](#) (page 11-39).

On platforms supporting named sockets, character device files or block device files, snapshot-based replication does not replicate these file types. However, the `acfsutil repl upgrade` process preserves any files of these types that already exist on the standby file system. They are not further updated, but are not deleted as part of the upgrade.

You must have system administrator privileges to run this command.

Example 16-36 Using the `acfsutil repl upgrade` command

```
# /sbin/acfsutil repl upgrade prepare -s my_repluser@my_host /my_primary/repl_data
```



```
# /sbin/acfsutil repl upgrade standby -u my_repluser /my_standby/repl_data
# /sbin/acfsutil repl upgrade primary -i 2h /my_primary/repl_data
```

16.8 Oracle ACFS Command-Line Tools for Auditing

This topic provides a summary of the commands for Oracle ACFS auditing.

[Table 16-41](#) (page 16-81) lists the Oracle ACFS auditing commands with brief descriptions. For an overview of Oracle ACFS auditing, refer to [Oracle ACFS Auditing](#) (page 11-29).

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Table 16-41 Summary of commands for Oracle ACFS auditing

Command	Description
acfsutil audit archive (page 16-81)	Forces an archival of the audit trail.
acfsutil audit disable (page 16-82)	Disables auditing for Oracle ACFS encryption or security.
acfsutil audit enable (page 16-82)	Enables auditing for Oracle ACFS encryption or security.
acfsutil audit info (page 16-83)	Displays auditing information.
acfsutil audit init (page 16-84)	Initializes auditing on an Oracle ACFS environment.
acfsutil audit purge (page 16-85)	Purges the audit trail for a specified file system.
acfsutil audit read (page 16-85)	Marks the audit trail to indicate that the audit trail has been read.

16.8.1 acfsutil audit archive

Purpose

Forces an archival of the audit trail for the current host on the specified Oracle ACFS file system.

Syntax and Description

```
acfsutil audit archive -h
acfsutil audit archive -m mount_point
```

`acfsutil audit archive -h` displays help text and exits.

[Table 16-42](#) (page 16-81) contains the options available with the `acfsutil audit archive` command.

Table 16-42 Options for the acfsutil audit archive command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.

You can use the `acfsutil audit archive` command to safely back up or remove audit data without the possibility of losing intermediate audit messages.

Only an audit manager can run this command.

Examples

The following example shows the use of the `acfsutil audit archive` command.

Example 16-37 Using the acfsutil audit archive command

```
$ /sbin/acfsutil audit archive -m /acfsmounts/acfs1
```

16.8.2 acfsutil audit disable**Purpose**

Disables auditing for either Oracle ACFS encryption or security on a specified file system.

Syntax and Description

```
acfsutil audit disable -h
acfsutil audit disable -m mount_point -s {encr |sec}
```

`acfsutil audit disable -h` displays help text and exits.

[Table 16-43](#) (page 16-82) contains the options available with the `acfsutil audit disable` command.

Table 16-43 Options for the acfsutil audit disable command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-s {encr sec}</code>	Specifies whether to disable auditing for encryption or security.

Only an audit manager can run this command.

Examples

The following example shows the use of the `acfsutil audit disable` command.

Example 16-38 Using the acfsutil audit disable command

```
$ /sbin/acfsutil audit disable -m /acfsmounts/acfs1 -s encr
```

16.8.3 acfsutil audit enable**Purpose**

Enables auditing for either Oracle ACFS encryption or security on a specified file system.

Syntax and Description

```
acfsutil audit enable -h
acfsutil audit enable -m mount_point -s {encr |sec}
```

`acfsutil audit enable -h` displays help text and exits.

[Table 16-44](#) (page 16-83) contains the options available with the `acfsutil audit enable` command.

Table 16-44 Options for the `acfsutil audit enable` command

Option	Description
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-s {encr sec}</code>	Specifies whether to enable auditing for encryption or security.

The `acfsutil audit enable` command may also create the `mount_point/.Security/audit` directory, which is a location for audit source files. The audit directory is created when auditing first is enabled for a file system, as a result of either the `acfsutil encr set`, `acfsutil sec prepare`, or `acfsutil audit enable` command.

Only an audit manager can run this command.

Examples

The following example shows the use of the `acfsutil audit enable` command.

Example 16-39 Using the `acfsutil audit enable` command

```
$ /sbin/acfsutil audit enable -m /acfsmounts/acfs1 -s encr
```

16.8.4 `acfsutil audit info`

Purpose

Displays auditing information.

Syntax and Description

```
acfsutil audit info -h
acfsutil audit info [-m mount_point ]
```

`acfsutil audit info -h` displays help text and exits.

[Table 16-45](#) (page 16-83) contains the options available with the `acfsutil audit info` command.

Table 16-45 Options for the `acfsutil audit info` command

Option	Description
<code>-m mount_point</code>	Optionally specifies the directory where the file system is mounted.

Only an audit manager can run this command.

Examples

The following example shows the use of the `acfsutil audit info` command.

Example 16-40 Using the `acfsutil audit info` command

```
$ /sbin/acfsutil audit info -m /acfsmounts/acfs1
```

```
Auditing information for '/acfsmounts/acfs1':
Audit trail size: 10MB
```

```
Archive File: READ
Audit Sources:
Security: ENABLED
Encryption: ENABLED
```

16.8.5 acfsutil audit init

Purpose

Initializes Oracle ACFS auditing.

Syntax and Description

```
acfsutil audit init -h
acfsutil audit init -M audit_manager_group -A auditor_group
```

`acfsutil audit init -h` displays help text and exits.

[Table 16-46](#) (page 16-84) contains the options available with the `acfsutil audit init` command.

Table 16-46 Options for the `acfsutil audit init` command

Option	Description
<code>-M <i>audit_manager_group</i></code>	Specifies the operating system (OS) group which assigns users to the Oracle ACFS audit manager role. To achieve separation of duties, you should create a new group specifically for this purpose and this group should be different from the system administrator group, Oracle ACFS security administrator OS group, and the Oracle ACFS audit auditor OS group.
<code>-A <i>auditor_group</i></code>	Specifies the operating system (OS) group which assigns users to the Oracle ACFS audit auditor role. To achieve separation of duties, you should create a new group specifically for this purpose and this group should be different from the system administrator group, Oracle ACFS security administrator OS group, and the Oracle ACFS audit manager OS group.

The `acfsutil audit init` command must be run by the system administrator before enabling auditing for any of the audit sources on a file system. The command sets up the required roles for auditing and must be run before any type of auditing can be enabled on a file system. After running the command, auditing is enabled by default for all features which use auditing.

Because you cannot choose a different OS group for either the Oracle ACFS audit manager or auditor after initialization, you should create new OS groups for these specific purposes before initializing Oracle ACFS auditing.

Only a system administrator can run this command.

Examples

The following example shows the use of the `acfsutil audit init` command.

Example 16-41 Using the `acfsutil audit init` command

```
# /sbin/acfsutil audit init -M myaudit_mgr_grp -A myauditor_grp
```

16.8.6 acfsutil audit purge

Purpose

Purges the audit trail for a specified file system.

Syntax and Description

```
acfsutil audit purge -h
acfsutil audit purge -m mount_point [-f]
```

`acfsutil audit purge -h` displays help text and exits.

[Table 16-47](#) (page 16-85) contains the options available with the `acfsutil audit purge` command.

Table 16-47 Options for the `acfsutil audit purge` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-f</code>	Forces the removal of the audit trail even if the audit auditor has not marked the file as read using the <code>acfsutil audit read</code> command.

Only an audit manager can run this command.

Examples

The following example shows the use of the `acfsutil audit purge` command.

Example 16-42 Using the `acfsutil audit purge` command

```
$ /sbin/acfsutil audit purge -m /acfsmounts/acfs1 -f
```

16.8.7 acfsutil audit read

Purpose

Marks the audit trail to indicate to the audit manager that the log archive file for the current node has been reviewed, backed up as necessary, and is safe to purge.

Syntax and Description

```
acfsutil audit read -h
acfsutil audit read -m mount_point
```

`acfsutil audit read -h` displays help text and exits.

[Table 16-48](#) (page 16-85) contains the options available with the `acfsutil audit read` command.

Table 16-48 Options for the `acfsutil audit read` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.

Only an audit auditor can run this command.

Examples

The following example shows the use of the `acfsutil audit read` command.

Example 16-43 Using the `acfsutil audit read` command

```
$ /sbin/acfsutil audit read -m /acfsmounts/acfs1
```

16.9 Oracle ACFS Command-Line Tools for Security

This topic provides a summary of the commands for Oracle ACFS security.

[Table 16-49](#) (page 16-86) lists the Oracle ACFS security commands with brief descriptions. For an overview of Oracle ACFS security, refer to [Oracle ACFS Security](#) (page 11-22).

For more information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Table 16-49 Summary of commands for Oracle ACFS security

Command	Description
acfsutil sec admin add (page 16-87)	Adds a security administrator.
acfsutil sec admin password (page 16-88)	Changes the password of a security administrator.
acfsutil sec admin remove (page 16-88)	Removes a security administrator.
acfsutil sec batch (page 16-89)	Runs a batch file.
acfsutil sec disable (page 16-90)	Disables Oracle ACFS security.
acfsutil sec enable (page 16-91)	Enables Oracle ACFS security.
acfsutil sec info (page 16-91)	Displays Oracle ACFS file system security information.
acfsutil sec info file (page 16-92)	Lists the security realms that a specified file or directory belongs to.
acfsutil sec init (page 16-93)	Initializes Oracle ACFS file system security.
acfsutil sec load (page 16-95)	Loads Oracle ACFS file system security metadata.
acfsutil sec prepare (page 16-96)	Prepares an Oracle ACFS file system for security.
acfsutil sec realm add (page 16-98)	Adds objects to an Oracle ACFS file system security realm.
acfsutil sec realm audit disable (page 16-102)	Disables auditing of command rules for files in an Oracle ACFS security realm.
acfsutil sec realm audit enable (page 16-103)	Enables auditing of command rules for files in an Oracle ACFS security realm.
acfsutil sec realm audit info (page 16-105)	Displays the realm auditing information for a specified Oracle ACFS security realm.
acfsutil sec realm clone (page 16-105)	Clones an Oracle ACFS file system security realm.
acfsutil sec realm create (page 16-107)	Creates an Oracle ACFS file system security realm.

Table 16-49 (Cont.) Summary of commands for Oracle ACFS security

Command	Description
acfsutil sec realm delete (page 16-108)	Removes objects from an Oracle ACFS file system security realm.
acfsutil sec realm destroy (page 16-109)	Removes an Oracle ACFS file system security realm.
acfsutil sec rule clone (page 16-109)	Clones an Oracle ACFS file system security rule.
acfsutil sec rule create (page 16-110)	Creates an Oracle ACFS file system security rule.
acfsutil sec rule destroy (page 16-112)	Removes an Oracle ACFS file system security rule.
acfsutil sec rule edit (page 16-112)	Updates an Oracle ACFS file system security rule.
acfsutil sec ruleset clone (page 16-113)	Clones an Oracle ACFS file system security rule set.
acfsutil sec ruleset create (page 16-114)	Creates an Oracle ACFS file system security rule set.
acfsutil sec ruleset destroy (page 16-115)	Removes an Oracle ACFS file system rule set.
acfsutil sec ruleset edit (page 16-116)	Updates an Oracle ACFS file system rule set.
acfsutil sec save (page 16-116)	Saves Oracle ACFS file system security metadata.

16.9.1 acfsutil sec admin add

Purpose

Adds a new security administrator for an Oracle ACFS file system.

Syntax and Description

```
acfsutil sec admin add -h
acfsutil sec admin add admin
```

`acfsutil sec admin add -h` displays help text and exits.

[Table 16-50](#) (page 16-87) contains the options available with the `acfsutil sec admin add` command.

Table 16-50 Options for the acfsutil sec admin add command

Option	Description
<i>admin</i>	Specifies a security administrator user name. The user specified must be an existing operating system user and a member of the security group specified with the <code>acfsutil sec init</code> command. On Windows, a security administrator user name must be specified with a fully qualified domain user name in the form of <i>domain_name\username</i> .

Security administrators are common for all Oracle ACFS file systems in a cluster. A temporary password must be provided for the new security administrator. The password must conform to the format that is described in "[acfsutil sec init](#) (page 16-93)".

The new security administrator can change the password with the `acfsutil sec admin password` command. For information, refer to "[acfsutil sec admin password](#) (page 16-88)".

Security administrators are allowed to browse all directories in an Oracle ACFS file system whether they have the underlying operating system permissions and whether any realm checks allow it. This exception enables a security administrator to check the location of the files when securing them with Oracle ACFS security realms. However, a security administrator cannot view the contents of individual files without the appropriate operating system and security realm permissions.

Only an existing security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec admin add` command.

Example 16-44 Using the `acfsutil sec admin add` command

```
$ /sbin/acfsutil sec admin add sec_admin_three
```

16.9.2 `acfsutil sec admin password`

Purpose

Changes the password of a security administrator for an Oracle ACFS file system.

Syntax and Description

```
acfsutil sec admin password -h
acfsutil sec admin password
```

`acfsutil sec admin password -h` displays help text and exits.

The `acfsutil sec admin password` command changes the security password for the administrator that is running the command. When you run this command, you are prompted to enter a new password. The password must conform to the format that is described in "[acfsutil sec init](#) (page 16-93)".

Every time a security administrator runs an `acfsutil sec` command, the administrator is prompted for the security administrator's password.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec admin password` command.

Example 16-45 Using the `acfsutil sec admin password` command

```
$ /sbin/acfsutil sec admin password
ACFS Security administrator password:
New password:
Re-enter new password:
```

16.9.3 `acfsutil sec admin remove`

Purpose

Removes a security administrator from an Oracle ACFS file system.

Syntax and Description

```
acfsutil sec admin remove -h
acfsutil sec admin remove admin
```

`acfsutil sec admin remove -h` displays help text and exits.

[Table 16-51](#) (page 16-89) contains the options available with the `acfsutil sec admin remove` command.

Table 16-51 Options for the `acfsutil sec admin remove` command

Option	Description
<i>admin</i>	Specifies an existing security administrator user name. On Windows, the security administrator user name must be specified with a fully qualified user name in the form of <i>domain_name\username</i> .

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec admin remove` command.

Example 16-46 Using the `acfsutil sec admin remove` command

```
$ /sbin/acfsutil sec admin remove sec_admin_three
```

16.9.4 acfsutil sec batch

Purpose

Runs a specified batch file.

Syntax and Description

```
acfsutil sec batch -h
acfsutil sec batch batch_file
```

`acfsutil sec batch -h` displays help text and exits.

[Table 16-52](#) (page 16-89) contains the options available with the `acfsutil sec batch` command.

Table 16-52 Options for the `acfsutil sec batch` command

Option	Description
<i>batch_file</i>	Specifies an existing batch file name. The batch file contains a list of <code>acfsutil sec</code> commands.

The batch file can only contain security realm management commands. Interactive commands are not recommended. The `acfsutil sec admin add`, `acfsutil sec admin password`, and `acfsutil sec init` commands are not supported in the batch file. Also, other `acfsutil` commands, such as `acfsutil encr` commands, are not allowed in the batch file. If a command in the batch file fails, subsequent commands in the batch file are not run.

The following are examples of commands that can be in a batch file:

```
acfsutil sec realm create my_realm1 -m /mnt1 -e off
acfsutil sec realm create my_realm2 -m /mnt2 -e off
```

Only a security administrator can run this command. When the command is run, the administrator is prompted once for a password.

Examples

The following example shows the use of the `acfsutil sec batch` command.

Example 16-47 Using the `acfsutil sec batch` command

```
$ /sbin/acfsutil sec batch my_batch_file
```

16.9.5 `acfsutil sec disable`

Purpose

Disables Oracle ACFS security on a mount point or a realm in a mount point.

Syntax and Description

```
acfsutil sec disable -h
acfsutil sec disable -m mount_point [-S snap_name] [realm]
```

`acfsutil sec disable -h` displays help text and exits.

[Table 16-53](#) (page 16-90) contains the options available with the `acfsutil sec disable` command.

Table 16-53 Options for the `acfsutil sec disable` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-S <i>snap_name</i></code>	Disables security for the specified read-write snapshot.
<code><i>realm</i></code>	Specifies the name of the security realm in the Oracle ACFS file system.

The `acfsutil sec disable -m mount_point` command disables security functionality on the Oracle ACFS file system specified by the mount point option. When security is disabled on the file system, security realms do not enforce realm authorization.

The `acfsutil sec disable -m mount_point realm` command disables security for the realm specified in the command.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec disable` command.

Example 16-48 Using the `acfsutil sec disable` command

```
$ /sbin/acfsutil sec disable -m /acfsmounts/acfs1 my_realm
```

16.9.6 acfsutil sec enable

Purpose

Enables Oracle ACFS security on a mount point or a realm in a mount point.

Syntax and Description

```
acfsutil sec enable -h
acfsutil sec enable -m mount_point [-S snap_name] [realm]
```

`acfsutil sec enable -h` displays help text and exits.

[Table 16-54](#) (page 16-91) contains the options available with the `acfsutil sec enable` command.

Table 16-54 Options for the `acfsutil sec enable` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-S <i>snap_name</i></code>	Enables security for the specified read-write snapshot.
<code><i>realm</i></code>	Specifies the name of the security realm.

The `acfsutil sec enable -m mount_point` command enables security functionality on the Oracle ACFS file system specified by the mount point option. When security is enabled on the file system, security realms that have been enabled enforce realm authorization. You should run this command before enabling any individual security realm.

The `acfsutil sec enable -m mount_point realm` command enables security for the realm specified in the command. The realm enforces authorization if security has been enabled on the file system.

Only a security administrator can run this command.

Examples

These example shows the use of the `acfsutil sec enable` command.

Example 16-49 Using the `acfsutil sec enable` command

```
$ /sbin/acfsutil sec enable -m /acfsmounts/acfs1
$ /sbin/acfsutil sec enable -m /acfsmounts/acfs1 my_realm
```

16.9.7 acfsutil sec info

Purpose

Displays information about Oracle ACFS security.

Syntax and Description

```
acfsutil sec info -h
acfsutil sec info -m mount_point
    [{-n [realm] | -l [rule] | -s [ruleset] | -c } ] [-S snap_name]
```

`acfsutil sec info -h` displays help text and exits.

Table 16-55 (page 16-92) contains the options available with the `acfsutil sec info` command.

Table 16-55 Options for the `acfsutil sec info` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-n <i>realm</i></code>	Displays information about the specified security realm. If the realm name is omitted, a list of all realms is displayed.
<code>-l <i>rule</i></code>	Displays information about the specified rule. If the rule name is omitted, a list of all rules is displayed.
<code>-s <i>ruleset</i></code>	Displays information about the specified rule set. If the rule set name is omitted, a list of all rule sets is displayed.
<code>-c</code>	Lists all the command rules.
<code>-S <i>snap_name</i></code>	Displays information about the realms, rules, and rule sets in the specified snapshot.

The `acfsutil sec info` command retrieves information about the list of realms, rules, and rule sets on the specified mount point. By specifying a particular realm, rule, or ruleset, you can retrieve information specific about the specified realm, rule, or ruleset. You can also display information about a specified snapshot.

If the `-m` option is specified without any other options, then the security enabled status and prepared status are displayed for the specified mount point.

To access files in the system security realms, the user should be assigned as a security administrator with the `acfsutil sec admin add` command. Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec info` command.

Example 16-50 Using the `acfsutil sec info` command

```
$ /sbin/acfsutil sec info -m /acfsmounts/acfs1 -n my_realm
```

16.9.8 acfsutil sec info file

Purpose

Lists the names of the Oracle ACFS security realms that the specified file or directory belongs to.

Syntax and Description

```
acfsutil sec info file -h
acfsutil sec info file -m mount_point path
```

`acfsutil sec info file -h` displays help text and exits.

[Table 16-56](#) (page 16-93) contains the options available with the `acfsutil sec info file` command.

Table 16-56 Options for the `acfsutil sec info file` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code><i>path</i></code>	Specifies the path of the file or directory in the file system.

This command also displays the encryption status of files.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec info file` command.

Example 16-51 Using the `acfsutil sec info file` command

```
$ /sbin/acfsutil sec info file -m /acfsmounts/acfs1
                               /acfsmounts/acfs1/myfiles
```

16.9.9 acfsutil sec init

Purpose

Initializes Oracle ACFS security.

Syntax and Description

```
acfsutil sec init -h
acfsutil sec init -u admin -g admin_sec_group
```

`acfsutil sec init -h` displays help text and exits.

[Table 16-57](#) (page 16-93) contains the options available with the `acfsutil sec init` command.

Table 16-57 Options for the `acfsutil sec init` command

Option	Description
<code>-u <i>admin</i></code>	Specifies the first security administrator user name. The user specified must be an existing operating system (OS) user and a member of the operating system group specified by the <code>-g</code> option. On Windows, the security administrator user name must be specified with a fully qualified user name in the form of <code>domain_name\username</code> .

Table 16-57 (Cont.) Options for the `acfsutil sec init` command

Option	Description
<code>-g admin_sec_group</code>	Specifies the name of the security group for the administrator. The group specified must be an existing operating system (OS) group. On Windows, the group name must be specified with a fully qualified domain group name in the form of <code>domain_name\groupname</code> . If the <code>domain_name\groupname</code> contains a space, then enclose the string in quotes (" ").

The `acfsutil sec init` command creates the storage necessary for security credentials and identifies an operating system user as the first security administrator. The command also identifies the operating system group that is the designated security group. All users that are security administrators must be members of the designated security group. Security administrators are common for all Oracle ACFS file systems.

If you are setting up an OS user and OS group, refer to your operating system-specific (OS) documentation for information.

The `acfsutil sec init` command is run once to set up Oracle ACFS security for each cluster and can be run from any node in the cluster. Other security commands can also be run from any node in a cluster.

Only the root user or Windows Administrator user can run this command. The user specifies a password for the security administrator. The security administrator password must conform to the following format:

- The maximum number of characters is 20.
- The minimum number of characters is 8.
- The password must contain at least one digit.
- The password must contain at least one letter.

The new security administrator can change the password with the `acfsutil sec admin password` command. For information, refer to "[acfsutil sec admin password](#) (page 16-88)".

Security administrators are allowed to browse all directories in an Oracle ACFS file system whether they have the underlying operating system permissions and whether any realm checks allow it. This exception enables a security administrator to check the location of the files when securing them with Oracle ACFS security realms. However, a security administrator cannot view the contents of individual files without the appropriate operating system and security realm permissions.

Examples

The following example shows the use of the `acfsutil sec init` command.

Example 16-52 Using the `acfsutil sec init` command

```
$ /sbin/acfsutil sec init -u grid -g asmadmin
```

16.9.10 acfsutil sec load

Purpose

Loads Oracle ACFS security metadata into a file system identified by a mount point.

Syntax and Description

```
acfsutil sec load -h
acfsutil sec load -m mount_point -p file
```

`acfsutil sec load -h` displays help text and exits.

[Table 16-58](#) (page 16-95) contains the options available with the `acfsutil sec load` command.

Table 16-58 Options for the `acfsutil sec load` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-p <i>file</i></code>	Specifies the name of an existing saved security metadata file.

The `acfsutil sec load` command loads the security metadata in a saved XML file into the specified Oracle ACFS file system. `acfsutil sec load` restores only user-created security policies; the command does not add files to the realms.

`acfsutil sec load` and `acfsutil sec save` can be used together to copy user-created policies from one file system to another. For example, if you have security policies on one file system that you want to replicate on other file systems, then use `acfsutil sec save` on the source file system to create an XML backup file. Next, use `acfsutil sec load` on the other destination file systems to load the saved security metadata and create the same policies. After creating the policies, you can choose to apply policies on different directories and files on that file system by adding directories and files in different realms, according to the policies you want to impose on those files.

To run the `acfsutil sec load` command, the destination mount point must have a file system that has been prepared for security and does not contain any user-created security objects.

If the file system mounted on destination mount point contains security objects, then you must run `acfsutil sec prepare -u` to remove all previously created security objects on the file system. After successfully running `acfsutil sec prepare -u`, you must run `acfsutil sec prepare` to prepare the file system for security. After successfully running `acfsutil sec prepare`, you can run `acfsutil sec load` on the file system. For information about preparing security on or removing security from a file system, refer to "[acfsutil sec prepare](#) (page 16-96)".

The `acfsutil sec load` command does not load system security realms from the backup file. System security realms are created with the `acfsutil sec prepare` command; `acfsutil sec load` does not re-create these realms. For information about the system-created security realms, refer to "[acfsutil sec prepare](#) (page 16-96)".

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec load` command.

Example 16-53 Using the `acfsutil sec load` command

```
$ /sbin/acfsutil sec load -m /acfsmounts/acfs1 -p my_metadata_file.xml
```

16.9.11 `acfsutil sec prepare`

Purpose

Prepares an Oracle ACFS file system for security features.

Syntax and Description

```
acfsutil sec prepare -h
acfsutil sec prepare [-u] -m mount_point
```

`acfsutil sec prepare -h` displays help text and exits.

[Table 16-59](#) (page 16-96) contains the options available with the `acfsutil sec prepare` command.

Table 16-59 Options for the `acfsutil sec prepare` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-u</code>	<p>Backs out security for the specified mount point. This command removes security from in the file system and reverts the file system to the state before <code>acfsutil sec prepare</code> was run on the file system.</p> <p>This command removes all realm-secured files and directories from the realms and then destroys all Oracle ACFS security rules, rule sets and realms from the file system. However, the <code>.Security</code> directory and its contents, including log files and the security metadata backup files, are not deleted.</p> <p>If you want to remove encryption and security is being used, then this command must be run before encryption is backed out. To back out encryption, refer to "acfsutil encr set (page 16-123)".</p>

The `acfsutil sec prepare` command must be run before any of the realm management commands. This command prepares the specified Oracle ACFS file system for security and by default turns security on for the file system.

When running `acfsutil sec prepare -u`, ensure that no other Oracle ACFS security commands are run until `acfsutil sec prepare` has completed.

If auditing is initialized on a cluster, this command also enables an Oracle ACFS security auditing source on the file system. The actions performed when enabling this audit source are the same as those done when the `acfsutil audit enable` command is run directly. For more information, refer to "[acfsutil audit enable](#) (page 16-82)".

This command creates the `/mount_point/.Security/`, `mount_point/.Security/backup`, and `/mount_point/.Security/realm/logs` directories where `mount_point` is the option specified in the command line.

This command creates the following system security realms:

- `SYSTEM_Logs`

This is a system-created realm to protect the Oracle ACFS security log files in the directory `.Security/realm/logs/` directory.

- `SYSTEM_Audit`

This is a system-created realm to protect audit trail files. This realm is created if auditing has been initialized. If auditing has not been initialized, it is created when auditing is enabled for the security source through the `acfsutil audit enable` command. This realm secures the audit trail file so that the audit manager can read and write and the auditor can read the file, and no one else has access. This realm also protects the audit trail file so the audit manager cannot delete (without running the `acfsutil audit purge` command), truncate, overwrite, or `chmod` the file.

- `SYSTEM_SecurityMetadata`

This is a system-created realm to protect the Oracle ACFS metadata XML file in the directory `.Security/backup/` directory.

- `SYSTEM_Antivirus`

This is a system-created realm that allows access for the antivirus software that is running on an Oracle ACFS file system. For every realm protected file or directory, the `SYSTEM_Antivirus` realm is evaluated when authorization checks are performed to determine if the `SYSTEM_Antivirus` realm allows access to the file or directory.

To allow the antivirus process to access realm-protected files or directories, you must add the `LocalSystem` or `SYSTEM` group to the realm with the `acfsutil sec realm add` command, as shown in [Example 16-55](#) (page 16-101). If other antivirus processes are running as `Administrator`, then the user `Administrator` must be added to the `SYSTEM_Antivirus` realm to allow access to realm protected files and directories.

If no Antivirus products have been installed, do not add any users or groups to the `SYSTEM_Antivirus` realm. Because users or groups added to the `SYSTEM_Antivirus` realm have `READ` and `READDIR` access, limit the users or groups added to this realm. You can restrict the time window when the users or groups of this realm can access the realm protected files or directories with time-based rules. You can also have application-based rules if you can identify the process name for the antivirus installation that scans the files.

The `SYSTEM_Antivirus` realm can only perform the following operations on a file or directory: `OPEN`, `READ`, `READDIR`, and setting time attributes. To remove or delete files or directories, you may need to disable security to clean up the infected files.

This realm is set up only for Windows systems.

- `SYSTEM_BackupOperators`

This is a system-created realm that enables you to authorize users that can back up realm-secured files and directories. You can add users, groups, rule sets, and command rules to this realm to provide fine-grain authorization for backing up realm-secured files and directories. A user must be added to this realm to back up realm-secured files and directories.

Use caution when adding groups to this system realm. After you add a group to this system realm, all the users of the added group are able to override the realm protections to access files.

To access files in the system security realms, the user should be assigned as a security administrator with the `acfsutil sec admin add` command.

You can add users, groups, rule sets, and command rules to system-created realms with the `acfsutil sec realm add` command, the same as for user-created realms. However, adding files and directories to system realms is not recommended. You can use the `acfsutil sec realm delete` command to delete objects from the system-created realms.

System-created security realms cannot be removed by a security administrator with the `acfsutil sec admin destroy` command. These realms are only removed when security is backed out of a file system when executing the `acfsutil sec prepare` command with the `-u` option.

The `acfsutil sec prepare -u` command is not allowed if any snapshots exist in the file system.

Only a security administrator can run the `acfsutil sec prepare` command.

Examples

The following example shows the use of the `acfsutil sec prepare` command.

Example 16-54 Using the `acfsutil sec prepare` command

```
$ /sbin/acfsutil sec prepare -m /acfsmounts/acfs1
```

16.9.12 `acfsutil sec realm add`

Purpose

Adds objects to an Oracle ACFS security realm.

Syntax and Description

```
acfsutil sec realm add -h
acfsutil sec realm add realm -m mount_point
    {[-u user, ...] [-G os_group,...]
    [-l commandrule:ruleset,commandrule:ruleset, ...]
    [-e [-a {AES}] [-k {128|192|256}]]
    [-f [-r] path ...]}
```

`acfsutil sec realm add -h` displays help text and exits.

[Table 16-60](#) (page 16-98) contains the options available with the `acfsutil sec realm add` command.

Table 16-60 Options for the `acfsutil sec realm add` command

Option	Description
<code>realm</code>	Specifies the realm name to add.
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-u user</code>	Specifies user names to add.

Table 16-60 (Cont.) Options for the `acfsutil sec realm add` command

Option	Description
<code>-G <i>os_group</i></code>	Specifies the operating system groups to add.
<code>-l <i>commandrule:ruleset</i></code>	Specifies the filters to add. The <code>commandrule</code> switch is used to add one or more command rules to the realm with a rule set. <code>ruleset</code> specifies the rule set associated with the command rule for this realm. Only one rule set can be included with each command rule. For a list of command rules, refer to Table 16-61 (page 16-100). To display a list of the command rules, use <code>acfsutil sec info</code> with the <code>-c</code> option. Refer to " acfsutil sec info (page 16-91)".
<code>-e</code>	Enables encryption on the realm. Turning encryption on for the realm causes all files contained in the realm to be encrypted. These files remain encrypted until they are no longer part of an encrypted realm. Files that are encrypted are not re-encrypted to match the new specified encryption parameters.
<code>-a {AES}</code>	Specifies the encryption algorithm for the realm.
<code>-k { 128 192 256 }</code>	Specifies the encryption key length.
<code>-f [-r] <i>path</i> ...</code>	Adds files specified by <code>path</code> to the realm. <code>-r</code> specifies a recursive operation. File paths must be separated by spaces and must be placed at the end of the command. If a specified file is not realm secured, the file is encrypted or decrypted to match the encryption status for the realm.

The `acfsutil sec realm add` command adds objects to the specified realm. The objects to be added include users, groups, command rules, rule sets, and files. If the command encounters an error when adding an object, a message is displayed and the command continues processing the remaining objects.

Multiple entries can be added in a comma-delimited list when adding users, operating system groups, or command rules. Do not use spaces in the comma-delimited list. If spaces are added, then enclose the list in quotes.

If the `-e` option is specified, then encryption must have been initialized for the cluster and set on the file system. For more information, refer to "[acfsutil encr init](#) (page 16-119)" and "[acfsutil encr set](#) (page 16-123)".

If the entire mount point, which includes the `.Security` directory, is added to the realm then the security administrator operating system group should be added to the realm to maintain security logging and backing up operations.

The supported command are rules listed in [Table 16-61](#) (page 16-100). These command rules restrict or protect against file system operations on realm-secured files and directories.

Table 16-61 Security Realm Command Rules

Rule	Description
ALL	Protects against all file system operations on files and directories.
APPENDFILE	Restricts against additions to the end of a file. Restrictions include writes that start within the current file size, but proceed beyond the end of the file.
CHGRP	Protects from changing the group ownership on a file or directory.
CHMOD	Protects from changing the permissions on a file or directory.
CHOWN	Protects from changing the owner information of a file or directory.
CREATEFILE	Protects from creation of new file in a directory.
DELETEFILE	Protects from deletion of a file from a directory.
EXTEND	Restricts the extension operation of a file size. A file size may still be modifiable with other operations. <code>EXTEND</code> does not protect against a <code>truncate</code> followed by an <code>append</code> operation.
IMMUTABLE	<p>Denies any changes to the files and directories in the realm except changes to extended attributes resulting from commands such as <code>acfsutil tag</code> and <code>acfsutil encr</code>.</p> <p>Includes the following protection for a file or directory: <code>APPENDFILE</code>, <code>CHGRP</code>, <code>CHMOD</code>, <code>CHOWN</code>, <code>DELETEFILE</code>, <code>EXTEND</code>, <code>OVERWRITE</code>, <code>RENAME</code>, <code>RMDIR</code>, <code>TRUNCATE</code>, and <code>WRITE</code>.</p> <p><code>IMMUTABLE</code> does not deny any changes to the <code>atime</code> attribute. The <code>atime</code> attribute changes when a user accesses the file.</p> <p>Can be set to archive the files and directories in a security realm.</p>
LINKFILE	Restricts the creation of hard links to files.
MKDIR	Protects from the creation of new directory in a directory.
MMAPREAD	Protects a file from being memory mapped for a read operation using <code>mmap()</code> on Linux or using <code>CreateFileMapping</code> followed by <code>MapViewOfFile()</code> on Windows.
MMAPWRITE	Protects a file from being memory mapped for a write operation. Setting <code>MMAPWRITE</code> also protects a file from mapping for read as the operating system maps a file for both read and write.
OPENFILE	Protects from the opening of a file.

Table 16-61 (Cont.) Security Realm Command Rules

Rule	Description
OVERWRITE	Prevents existing content in a file from being overwritten with a <code>write</code> operation whose start and end offsets are within the current file size. If the operations on a file are <code>truncate</code> followed by <code>append</code> , <code>OVERWRITE</code> does not protect the file. To provide additional protection from both <code>append</code> and <code>overwrite</code> operations, use the <code>WRITE</code> command rule.
REaddir	Restricts for a directory listing, except for use by the security administrator group.
READ	Protects from reading the contents of a file. <code>READ</code> also protects against <code>read</code> operations using <code>mmap(2)</code> .
RENAME	Protects against renaming a file or directory.
RMDIR	Protects against removing a directory.
SYMLINK	Restricts the creation of symbolic links in the directories protected by a security realm. When creating symbolic links, it does not matter whether the source file is protected by a security realm.
TRUNCATE	Restricts the truncation of a file.
WRITE	Protects a file against the <code>write</code> system call. <code>WRITE</code> also protects against <code>append</code> and <code>overwrite</code> operations, plus <code>write</code> operations using <code>mmap(2)</code> . A file may still be modifiable with other file operations. To protect the file from other modifications, also use the <code>TRUNCATE</code> and <code>DELETEFILE</code> command rules.

Only a security administrator can run this command.

Examples

[Example 16-55](#) (page 16-101) shows the use of the `acfsutil sec realm add` command. The first `acfsutil sec` command adds a user group to a security realm. The second and third commands add the `LocalSystem` or `SYSTEM` group to the `SYSTEM_Antivirus` realm in a Windows environment.

Example 16-55 Using the `acfsutil sec realm add` command

```
$ /sbin/acfsutil sec realm add my_security_realm -m /acfsmounts/acfs1
-G my_os_group

C:\> acfsutil sec realm add SYSTEM_Antivirus /m e: /G "NT AUTHORITY\SYSTEM"

C:\> acfsutil sec realm add SYSTEM_Antivirus /m e: /G "SYSTEM"
```

16.9.13 acfsutil sec realm audit disable

Purpose

Disables auditing of a specific command rule or all command rules for files in an Oracle ACFS security realm.

Syntax and Description

```
acfsutil sec realm audit disable -h
acfsutil sec realm audit disable realm -m mount_point
    [-l commandrule,commandrule,...] {-a |-v }
```

`acfsutil sec realm audit disable -h` displays help text and exits.

[Table 16-62](#) (page 16-102) contains the options available with the `acfsutil sec realm audit disable` command.

Table 16-62 Options for the `acfsutil sec realm audit disable` command

Option	Description
<i>realm</i>	Specifies the security realm name.
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-l commandrule</code>	Specifies the command rules on which to disable auditing. If this option is not specified, then the list of all command rules is the default. For a list of command rules, refer to Table 16-61 (page 16-100). To display a list of the command rules, use <code>acfsutil sec info</code> with the <code>-c</code> option. Refer to " acfsutil sec info (page 16-91)".
<code>-a -v</code>	Specifies to disable audit realm authorizations (<code>-a</code>) or disable audit realm violations (<code>-v</code>). Either <code>-a</code> or <code>-v</code> must be specified.

Multiple entries can be added in a comma-delimited list when listing command rules. Do not use spaces in the comma-delimited list. If spaces are added, then enclose the list in quotes.

Only a security administrator can run this command. This command is authenticated using the Oracle ACFS security administrator password.

Examples

[Example 16-56](#) (page 16-102) shows the use of the `acfsutil sec realm audit disable` command. This command disables auditing on the OPEN (all violations) and WRITE (all violations) command rules.

Example 16-56 Using the `acfsutil sec realm audit disable` command

```
$ /sbin/acfsutil sec realm audit disable mySecureRealm
    -m /acfsmounts/acfs1 -l OPEN,WRITE -v
```

16.9.14 acfsutil sec realm audit enable

Purpose

Enables auditing of a specific command rule or all command rules for files in an Oracle ACFS security realm.

Syntax and Description

```
acfsutil sec realm audit enable -h
acfsutil sec realm audit enable realm -m mount_point
    [-l commandrule,commandrule,...]
    [-a ] [-v [ -u] ]
```

`acfsutil sec realm audit enable -h` displays help text and exits.

[Table 16-63](#) (page 16-103) contains the options available with the `acfsutil sec realm audit enable` command.

Table 16-63 Options for the `acfsutil sec realm audit enable` command

Option	Description
<i>realm</i>	Specifies the security realm name.
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-l commandrule</code>	Specifies the command rules on which to enable auditing. If this option is not specified, then the list of all command rules is the default. For a list of command rules, refer to Table 16-61 (page 16-100). To display a list of the command rules, use <code>acfsutil sec info</code> with the <code>-c</code> option. Refer to " acfsutil sec info " (page 16-91)".
<code>-a</code>	Specifies to audit realm authorizations.
<code>-v [-u]</code>	Specifies to audit realm violations. If <code>-u</code> is also specified, only realm violations by users who are members of a realm are audited.

If the `acfsutil sec realm audit enable` command is run multiple times, then the earlier configuration is not negated and the new settings are also applied. An exception to this behavior occurs when the command is run with the `-v` option and the specified command rule has auditing set for realm violations. In this case, the behavior is updated according to whether the `-u` flag was specified. For more information, see [Example 16-59](#) (page 16-104).

Multiple entries can be added in a comma-delimited list when listing command rules. Do not use spaces in the comma-delimited list. If spaces are added, then enclose the list in quotes.

If neither `-a` or `-v` are specified with the `acfsutil sec realm audit enable` command, the default is `-v`. Both `-a` and `-v` can be specified.

Only a security administrator can run this command. This command is authenticated using the Oracle ACFS security administrator password.

Examples

[Example 16-57](#) (page 16-104) shows how to enable auditing of the Oracle ACFS backup operators. Because these users are allowed access to files through the `SYSTEM_Backup` realm and are granted special privileges that give them access to all files on the file system, a security administrator may want to audit their actions. After the command is executed, any time a member of the `SYSTEM_Backup` realm opens a file an audit record is written to the Oracle ACFS Security audit trail on the file system.

Example 16-57 Auditing Oracle ACFS security backup operators

```
$ /sbin/acfsutil sec realm audit enable SYSTEM_Backup
-m /acfsmounts/acfs1 -l OPEN -a
```

[Example 16-58](#) (page 16-104) shows how to use the `-u` option to audit realm violations by users who are part of the realm. In this scenario sensitive human resources information is stored in `HumanResources` security realm and the `hr` group is allowed to access this information. However, a ruleset applied to the `ALL` command rule prevents access to this data from 6 PM to 8 AM. With this command, the security administrator could discover if any human resource employees are attempting to access sensitive data outside of the allowed time period. After this command is executed, only access violations by users who are members of the `hr` group are audited.

Example 16-58 Auditing only security realm users

```
$ /sbin/acfsutil sec realm audit enable HumanResources
-m /acfsmounts/acfs1 -l ALL -v -u
```

[Example 16-59](#) (page 16-104) shows multiple runs of the `acfsutil sec realm audit enable` command. After run 1, the `OPEN` (all violations) and `WRITE` (all violations) command rules are audited. After run 2, the `OPEN` (all violations), `WRITE` (all violations), and `DELETEFILE` (authorizations) command rules are audited. After run 3, the `OPEN` (authorizations and realm user violations), `WRITE` (all violations), `DELETEFILE` (authorizations), and `TRUNCATE` (authorizations and realm user violations) command rules are audited. After run 4, all violations are audited on all command rules. In addition, authorizations are audited for `OPEN`, `DELETEFILE`, and `TRUNCATE`.

Example 16-59 Running acfsutil sec realm audit enable multiple times

```
$ echo run 1
$ /sbin/acfsutil sec realm audit enable mySecureRealm
-m /acfsmounts/acfs1 -l OPEN,WRITE -v

$ echo run 2
$ /sbin/acfsutil sec realm audit enable mySecureRealm
-m /acfsmounts/acfs1 -l DELETEFILE -a

$ echo run 3
$ /sbin/acfsutil sec realm audit enable mySecureRealm
-m /acfsmounts/acfs1 -l OPEN,TRUNCATE -a -v -u

$ echo run 4
$ /sbin/acfsutil sec realm audit enable mySecureRealm
-m /acfsmounts/acfs1 -v
```


16.9.15 acfsutil sec realm audit info

Purpose

Displays the realm auditing information for a specified Oracle ACFS security realm.

Syntax and Description

```
acfsutil sec realm audit info -h
acfsutil sec realm audit info -m mount_point -n realm
```

`acfsutil sec realm audit info -h` displays help text and exits.

[Table 16-63](#) (page 16-103) contains the options available with the `acfsutil sec realm audit info` command.

Table 16-64 Options for the `acfsutil sec realm audit info` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-n <i>realm</i></code>	Specifies the security realm name.

The `acfsutil sec realm audit info` command provides information about a specified Oracle ACFS security realm.

Examples

[Example 16-60](#) (page 16-105) shows an example of the `acfsutil sec realm audit info` command.

Example 16-60 Running `acfsutil sec realm audit info`

```
$ /sbin/acfsutil sec realm audit info -m /acfsmounts/acfs1
-n mySecureRealm

Command rule auditing information for realm 'mySecureRealm'
on mount point '/acfsmounts/acfs1':

Realm authorization           : 'READ, WRITE'
Realm violation for all users : 'READ, OPENFILE'
Realm violation for realm users: 'None'
```

16.9.16 acfsutil sec realm clone

Purpose

Clones an Oracle ACFS security realm.

Syntax and Description

```
acfsutil sec realm clone -h
acfsutil sec realm clone realm -s src_mount_point new_realm
[-e] [-f] [-G] [-l] [-u]
acfsutil sec realm clone realm -s src_mount_point
[new_realm] -d destination_mount_point
[-e] [-G] [-l] [-u]
```

`acfsutil sec realm clone -h` displays help text and exits.

Table 16-65 (page 16-106) contains the options available with the `acfsutil sec realm clone` command.

Table 16-65 Options for the `acfsutil sec realm clone` command

Option	Description
<code>realm</code>	Specifies the realm name to be cloned.
<code>-s src_mount_point</code>	Specifies the directory where the source file system is mounted.
<code>new_realm</code>	Specifies the new realm name.
<code>-d destination_mount_point</code>	Specifies the directory for the destination mount point for the new realm.
<code>-e</code>	Copy encryption attributes to the new realm.
<code>-f</code>	Copy file objects to the new realm.
<code>-G</code>	Copy operating system groups to the new realm.
<code>-l</code>	Copy filters to the new realm.
<code>-u</code>	Copy users to the new realm.

The `acfsutil sec realm clone` makes a copy of the specified realm in the destination mount point. If the source and mount points are different and the new realm name is not specified, then the realm is cloned using the existing realm name in the Oracle ACFS file system specified by destination mount point. If the destination mount point is not specified, then the cloned realm is located in the source mount point and a new unique realm name must be specified.

If the `-l` option is specified and the destination mount point is different than the source mount point, then the rules and rule sets must be cloned first.

If the `-e` option is specified and the destination mount point is different than the source mount point, then encryption must be set on destination mount point. For more information, refer to "[acfsutil encr set](#) (page 16-123)".

The `-f` option can only be used if the destination mount point is the same as the source mount point.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec realm clone` command.

Example 16-61 Using the `acfsutil sec realm clone` command

```
$ /sbin/acfsutil sec realm clone my_security_realm -s /acfsmounts/acfs1
my_new_security_realm -d /acfsmounts/acfs2 -G
```

16.9.17 acfsutil sec realm create

Purpose

Creates an Oracle ACFS security realm.

Syntax and Description

```
acfsutil sec realm create -h
acfsutil sec realm create realm -m mount_point
-e { on -a {AES} -k {128|192|256} | off }
[-o {enable|disable}] [-d "description"]
```

`acfsutil sec realm create -h` displays help text and exits.

[Table 16-66](#) (page 16-107) contains the options available with the `acfsutil sec realm create` command.

Table 16-66 Options for the `acfsutil sec realm create` command

Option	Description
<i>realm</i>	Specifies the realm name.
<code>-m mount_point</code>	Specifies the mount point for the file system. A mount point is specified as a path on Linux platforms.
<code>-e {on off}</code>	Specifies encryption on or off for the realm.
<code>-a {AES}</code>	Specifies the encryption algorithm.
<code>-k {128 192 256}</code>	Specifies the encryption key length.
<code>-o {enable disable}</code>	Specifies where security is on or off for the realm.
<code>-d "description"</code>	Specifies a realm description.

The `acfsutil sec create realm` creates a new realm in the specified Oracle ACFS file system. The new realm name must be unique in the file system identified by the mount point.

A maximum of 500 Oracle ACFS security realms can be created, including any default system realms created by the `acfsutil sec prepare` command.

The realm is enabled by default unless the `-o disable` option is specified.

If the `-e on` option is specified, then encryption must have been initialized for the cluster and set on the file system. For more information, refer to "[acfsutil encr init](#) (page 16-119)" and "[acfsutil encr set](#) (page 16-123)".

If the `-e off` option is specified, you cannot specify the `-a` and `-k` options.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec realm create` command.

Example 16-62 Using the `acfsutil sec realm create` command

```
$ /sbin/acfsutil sec realm create my_security_realm -m /acfsmounts/acfs1
-e on -a AES -k 192 -o enable
```

16.9.18 acfsutil sec realm delete**Purpose**

Deletes objects from an Oracle ACFS security realm.

Syntax and Description

```
acfsutil sec realm delete -h
acfsutil sec realm delete realm -m mount_point
    {[-u user, ...] [-G os_group, ...]
    [-l :ruleset,commandrule:ruleset, ...]
    [-f [-r] path, ...] } [-e ]}
```

`acfsutil sec realm delete -h` displays help text and exits.

[Table 16-67](#) (page 16-108) contains the options available with the `acfsutil sec realm delete` command.

Table 16-67 Options for the `acfsutil sec realm delete` command

Option	Description
<i>realm</i>	Specifies the realm name.
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-u user</code>	Specifies user names to delete.
<code>-G os_group</code>	Specifies the operating system groups to delete.
<code>-l commandrule:ruleset</code>	Specifies the filters to delete from the realm. To display a list of the command rules, use <code>acfsutil sec info</code> with the <code>-c</code> option. <i>ruleset</i> specifies the rule set associated with the command rule for this realm.
<code>-f [-r] path ...</code>	Deletes files specified by <i>path</i> from the realm. <code>-r</code> specifies a recursive operation. File paths must be separated by spaces. If this is the last realm securing the file, the file is encrypted or decrypted to match the file system level encryption state.
<code>-e</code>	Disables encryption on the realm. When disabling encryption, this option decrypts any files in the realm that do not belong to any other encrypted realms. If a file is part of another realm which is encrypted or if encryption is turned on for the file system, then the file remains encrypted.

The `acfsutil sec realm delete` command removes objects from the specified realm. The objects to be deleted include users, groups, rule sets, and files. If the command encounters an error when deleting an object, a message is displayed and the command continues processing the remaining objects.

Multiple entries can be added in a comma-delimited list when adding users, operating system groups, or command rules. Do not use spaces in the comma-delimited list. If spaces are added, then enclose the list in quotes.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec realm delete` command.

Example 16-63 Using the `acfsutil sec realm delete` command

```
$ /sbin/acfsutil sec realm delete my_security_realm -m /acfsmounts/acfs1
-f -r /acfsmounts/acfs1/myoldfiles/*.log
```

16.9.19 `acfsutil sec realm destroy`

Purpose

Destroys an Oracle ACFS security realm.

Syntax and Description

```
acfsutil sec realm destroy -h
acfsutil sec realm destroy realm -m mount_point
```

`acfsutil sec realm destroy -h` displays help text and exits.

[Table 16-68](#) (page 16-109) contains the options available with the `acfsutil sec realm destroy` command.

Table 16-68 Options for the `acfsutil sec realm destroy` command

Option	Description
<i>realm</i>	Specifies the realm name.
<i>-m mount_point</i>	Specifies the directory where the file system is mounted.

The `acfsutil sec destroy realm` removes a security realm from the specified Oracle ACFS file system. Destroying the realm does not destroy the objects in the realm; this command simply removes the security associated with the realm from the objects.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec realm destroy` command.

Example 16-64 Using the `acfsutil sec realm destroy` command

```
$ /sbin/acfsutil sec realm destroy my_security_realm -m /acfsmounts/acfs1
```

16.9.20 `acfsutil sec rule clone`

Purpose

Clones a security rule.

Syntax and Description

```
acfsutil sec rule clone -h
acfsutil sec rule clone rule -s src_mount_point new_rule
acfsutil sec rule clone rule -s src_mount_point
    [new_rule] -d mount_point
```

`acfsutil sec rule clone -h` displays help text and exits.

[Table 16-69](#) (page 16-110) contains the options available with the `acfsutil sec rule clone` command.

Table 16-69 Options for the `acfsutil sec rule clone` command

Option	Description
<code>rule</code>	Specifies the existing name of the rule. If the name contains a space, enclose in quotes (" ").
<code>-s src_mount_point</code>	Specifies the directory where the source file system is mounted.
<code>-d mount_point</code>	Specifies the directory for the destination mount point of the file system.
<code>new_rule</code>	Specifies the new name of the rule. If the name contains a space, enclose in quotes (" ").

If the source and mount points are different and the new rule name is not specified, then the rule is cloned using the existing rule name in the Oracle ACFS file system specified by destination mount point. If the destination mount point is not specified, then the cloned rule is located in the source mount point and a new unique rule name must be specified.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec rule clone` command.

Example 16-65 Using the `acfsutil sec rule clone` command

```
$ /sbin/acfsutil sec rule clone my_security_rule -s /acfsmounts/acfs1
    my_new_security_rule -d /acfsmounts/acfs2
```

16.9.21 `acfsutil sec rule create`

Purpose

Creates a security rule.

Syntax and Description

```
acfsutil sec rule create -h
acfsutil sec rule create rule -m mount_point
    -t rule_type rule_value
    [-o {ALLOW|DENY}]
```

`acfsutil sec rule create -h` displays help text and exits.

[Table 16-70](#) (page 16-111) contains the options available with the `acfsutil sec rule create` command.

Table 16-70 Options for the `acfsutil sec rule create` command

Option	Description
<code>rule</code>	Specifies the name of the rule. If the name contains a space, enclose in quotes (" ").
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-t rule_type</code> <code>rule_value</code>	Specifies a rule type and a rule value. The rule type can be <code>application</code> , <code>hostname</code> , <code>time</code> , or <code>username</code> . The rule value depends on the type of rule. The valid rule types and values are described in this section.
<code>-o option</code>	Specifies options preceded by <code>-o</code> . The option specified can be <code>ALLOW</code> or <code>DENY</code> . The default value is <code>DENY</code> .

The `acfsutil sec rule create` command creates a new rule in the Oracle ACFS file system specified by the mount point. The new rule can be added to a rule set and that rule set can be added to a security realm.

A maximum of 500 Oracle ACFS security rules can be created.

The rule types and associated rule values are:

- `application`

This rule type specifies the name of an application which is allowed or denied access to the objects protected by a realm.

- `hostname`

This rule type specifies the name of a computer from which a user accesses the objects protected by a realm. Access from a node can be allowed or denied using this rule. The `hostname` should be one of the cluster node names and not any other external nodes which could have mounted the Oracle ACFS file system as a network File System (NFS) mount.

- `time`

This rule type specifies the time interval in the form `start_time,end_time`. This time interval specifies access to a realm. Access can be allowed or denied to objects protected by a realm only during certain times of the day by setting this rule in a realm. The time is based on the local time of the host.

- `username`

This rule type specifies the name of a user to be added or deleted from a realm. You can use this option to deny access for any user that belongs to a security group that is part of a realm.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec rule create` command.

Example 16-66 Using the `acfsutil sec rule create` command

```
$ /sbin/acfsutil sec rule create my_security_rule -m /acfsmounts/acfs1
-t username security_user_one -o ALLOW
```

16.9.22 acfsutil sec rule destroy

Purpose

Removes a security rule.

Syntax and Description

```
acfsutil sec rule destroy -h
acfsutil sec rule destroy rule -m mount_point
```

`acfsutil sec rule destroy -h` displays help text and exits.

[Table 16-71](#) (page 16-112) contains the options available with the `acfsutil sec rule destroy` command.

Table 16-71 Options for the `acfsutil sec rule destroy` command

Option	Description
<code>rule</code>	Specifies the name of the rule. If the name contains a space, enclose in quotes (" ").
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.

The `acfsutil sec rule destroy` command removes a rule from the rule sets in the Oracle ACFS file system specified by the mount point. A rule set is not destroyed if all the rules are destroyed. The empty rule set must be explicitly destroyed.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec rule destroy` command.

Example 16-67 Using the `acfsutil sec rule destroy` command

```
$ /sbin/acfsutil sec rule destroy my_security_rule -m /acfsmounts/acfs1
```

16.9.23 acfsutil sec rule edit

Purpose

Updates a security rule.

Syntax and Description

```
acfsutil sec rule edit -h
acfsutil sec rule edit rule -m mount_point
    { [-t rule_type rule_value ] [-o {ALLOW|DENY}] }
```

`acfsutil sec rule edit -h` displays help text and exits.

[Table 16-72](#) (page 16-113) contains the options available with the `acfsutil sec rule edit` command.

Table 16-72 Options for the `acfsutil sec rule edit` command

Option	Description
<code>rule</code>	Specifies the name of the rule. If the name contains a space, enclose in quotes (" ").
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-t rule_type</code> <code>rule_value</code>	Specifies a rule type and a rule value. The rule type can be <code>application</code> , <code>hostname</code> , <code>time</code> , or <code>username</code> . Rule value depends on the type of rule. For information on the rule type and rule value, refer to " acfsutil sec rule create (page 16-110)".
<code>-o option</code>	Specifies options preceded by <code>-o</code> . The option specified can be <code>ALLOW</code> or <code>DENY</code> .

The `acfsutil sec rule edit` updates a rule. The value that is associated with a rule can be updated, but not the rule type.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec rule edit` command to update `my_security_rule`. The existing rule is of type `username` and that value cannot be changed.

Example 16-68 Using the `acfsutil sec rule edit` command

```
$ /sbin/acfsutil sec rule edit my_security_rule -m /acfsmounts/acfs1
-t username security_user_three -o ALLOW
```

16.9.24 `acfsutil sec ruleset clone`

Purpose

Clones a security rule set.

Syntax and Description

```
acfsutil sec ruleset clone -h
acfsutil sec ruleset clone ruleset -s mount_point new_ruleset
acfsutil sec ruleset clone ruleset -s mount_point
[new_ruleset] -d mount_point
```

`acfsutil sec ruleset clone -h` displays help text and exits.

[Table 16-73](#) (page 16-113) contains the options available with the `acfsutil sec ruleset clone` command.

Table 16-73 Options for the `acfsutil sec ruleset clone` command

Option	Description
<code>rule_set</code>	Specifies the existing name of the rule set. If the name contains a space, enclose in quotes (" ").
<code>-s mount_point</code>	Specifies the directory where the source file system is mounted.

Table 16-73 (Cont.) Options for the `acfsutil sec ruleset clone` command

Option	Description
<code>-d mount_point</code>	Specifies the directory for the destination mount point of the file system.
<code>new_rule_set</code>	Specifies the new name of the rule set. If the name contains a space, enclose in quotes (" ").

If the source mount point is different from destination mount point, the rules in the rule set must be cloned first.

If the source and mount points are different and the new rule set name is not specified, then the rule set is cloned using the existing rule set name in the Oracle ACFS file system specified by destination mount point. If the destination mount point is not specified, then the cloned rule set is located in the source mount point and a new unique rule set name must be specified.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec ruleset clone` command.

Example 16-69 Using the `acfsutil sec ruleset clone` command

```
$ /sbin/acfsutil sec ruleset clone
   my_security_ruleset -s /acfsmounts/acfs1
   my_new_security_ruleset -d /acfsmounts/acfs2
```

16.9.25 `acfsutil sec ruleset create`

Purpose

Creates a security rule set.

Syntax and Description

```
acfsutil sec ruleset create -h
acfsutil sec ruleset create rule_set -m mount_point
    [-o {ALL_TRUE|ANY_TRUE}]
```

`acfsutil sec ruleset create -h` displays help text and exits.

[Table 16-74](#) (page 16-114) contains the options available with the `acfsutil sec ruleset create` command.

Table 16-74 Options for the `acfsutil sec ruleset create` command

Option	Description
<code>rule_set</code>	Specifies the name of the rule set. If the name contains a space, enclose in quotes (" ").
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.

Table 16-74 (Cont.) Options for the `acfsutil sec ruleset create` command

Option	Description
<code>-o option</code>	Specifies options preceded by <code>-o</code> . The option specified can be <code>ALL_TRUE</code> or <code>ANY_TRUE</code> . The default value is <code>ALL_TRUE</code> .

The `acfsutil sec ruleset create` command creates a new rule set in the specified mount point.

A maximum of 500 Oracle ACFS security rule sets can be created.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec ruleset create` command.

Example 16-70 Using the `acfsutil sec ruleset create` command

```
$ /sbin/acfsutil sec ruleset create
   my_security_ruleset -m /acfsmounts/acfs1 -o ANY_TRUE
```

16.9.26 `acfsutil sec ruleset destroy`

Purpose

Removes a security rule set.

Syntax and Description

```
acfsutil sec ruleset destroy -h
acfsutil sec ruleset destroy rule_set -m mount_point
```

`acfsutil sec ruleset destroy -h` displays help text and exits.

[Table 16-75](#) (page 16-115) contains the options available with the `acfsutil sec ruleset destroy` command.

Table 16-75 Options for the `acfsutil sec ruleset destroy` command

Option	Description
<code>rule_set</code>	Specifies the name of the rule set. If the name contains a space, enclose in quotes (" ").
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.

The `acfsutil sec ruleset destroy` command removes a rule set from the Oracle ACFS file system specified by the mount point. Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec ruleset destroy` command.

Example 16-71 Using the `acfsutil sec ruleset destroy` command

```
$ /sbin/acfsutil sec ruleset destroy
   my_security_ruleset -m /acfsmounts/acfs1
```

16.9.27 acfsutil sec ruleset edit**Purpose**

Updates a security rule set.

Syntax and Description

```
acfsutil sec ruleset edit -h
acfsutil sec ruleset edit rule_set -m mount_point
    { [-a rule,...] [-d rule,...] [-o {ALL_TRUE|ANY_TRUE}] }
```

`acfsutil sec ruleset edit -h` displays help text and exits.

[Table 16-76](#) (page 16-116) contains the options available with the `acfsutil sec ruleset edit` command.

Table 16-76 Options for the `acfsutil sec ruleset edit` command

Option	Description
<i>rule_set</i>	Specifies the name of the rule set. If the name contains a space, enclose in quotes (" ").
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-a rule</code>	Specifies the rule to add.
<code>-d rule</code>	Specifies the rule to remove.
<code>-o option</code>	Specifies options preceded by <code>-o</code> . The option specified can be <code>ALL_TRUE</code> or <code>ANY_TRUE</code> .

The `acfsutil sec ruleset edit` command updates a rule set in the Oracle ACFS file system specified by the mount point.

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec ruleset edit` command.

Example 16-72 Using the `acfsutil sec ruleset edit` command

```
$ /sbin/acfsutil sec ruleset edit
   my_security_ruleset -m /acfsmounts/acfs1
   -a my_new_rule -o ANY_TRUE
```

16.9.28 acfsutil sec save**Purpose**

Saves Oracle ACFS file system security metadata.

Syntax and Description

```
acfsutil sec save -h
acfsutil sec save -m mount_point -p file
```

`acfsutil sec save -h` displays help text and exits.

[Table 16-77](#) (page 16-117) contains the options available with the `acfsutil sec save` command.

Table 16-77 Options for the `acfsutil sec save` command

Option	Description
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.
<code>-p file</code>	Specifies a file name to store the security metadata. The file is saved in the <code>/mount_point/.Security/backup/</code> directory.

The `acfsutil sec save` command saves the security metadata for an Oracle ACFS file system to an XML file. By default, the file is saved in the `/mount_point/.Security/backup` directory.

This file can be backed up as a regular file by a backup application. System realms protect this file and allow only members of these realms to access this file and prevent all other users including the root user and system administrator from access. For information about the system-created security realms, refer to "[acfsutil sec prepare](#) (page 16-96)".

Only a security administrator can run this command.

Examples

The following example shows the use of the `acfsutil sec save` command.

Example 16-73 Using the `acfsutil sec save` command

```
$ /sbin/acfsutil sec save -m /acfsmounts/acfs1 -p my_metadata_file.xml
```

16.10 Oracle ACFS Command-Line Tools for Encryption

This topic provides a summary of the commands for Oracle ACFS encryption.

[Table 16-78](#) (page 16-117) lists the Oracle ACFS encryption commands with brief descriptions. For an overview of Oracle ACFS encryption, refer to [Oracle ACFS Encryption](#) (page 11-25).

The `acfsutil encr on`, `acfsutil encr off`, and `acfsutil encr rekey` commands are not supported on files greater than 128 megabytes. The operation of those commands on a file greater than 128 megabytes fails and the resulting error message displays alternative actions.

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Table 16-78 Summary of commands for Oracle ACFS encryption

Command	Description
acfsutil encr info (page 16-118)	Displays encryption-related information about Oracle ACFS file systems.

Table 16-78 (Cont.) Summary of commands for Oracle ACFS encryption

Command	Description
acfsutil encr init (page 16-119)	Creates storage for encryption keys.
acfsutil encr off (page 16-120)	Disables encryption for an Oracle ACFS file system.
acfsutil encr on (page 16-121)	Encrypts an Oracle ACFS file system.
acfsutil encr rekey (page 16-122)	Generates a new key and re-encrypts an Oracle ACFS file system.
acfsutil encr set (page 16-123)	Sets or changes encryption parameters for an Oracle ACFS file system.
acfsutil keystore migrate (page 16-125)	Migrates the encryption keystore.

16.10.1 acfsutil encr info

Purpose

Displays encryption-related information about Oracle ACFS file systems, directories, or files.

Syntax and Description

```
acfsutil encr info -h
acfsutil encr info -m mount_point [[-r] path [path ...]]
```

`acfsutil encr info -h` displays help text and exits.

[Table 16-79](#) (page 16-118) contains the options available with the `acfsutil encr info` command.

Table 16-79 Options for the acfsutil encr info command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-r</code>	Specifies recursive action under an existing directory folder identified by <i>path</i> .
<i>path</i>	Specifies the absolute or relative path of a directory. Multiple path values are allowed.

If `-m` is specified without a *path*, the encryption status, algorithm, and key length are displayed for the file system level.

If `-r` is specified with a *path*, the encryption status, algorithm, and key length are displayed for all objects under the directory specified by *path*.

The `acfsutil encr info` command displays encryption status and parameters for files in a snapshot if the files are specified with the *path* option.

This command fails when it is run on realm-secured objects.

Any user can run this command to display encryption information about a file system, directory, or file.

If the `acfsutil encr info` command is run as a system administrator, then the output displays the types of keystore used. The types are single-sign-on wallet in the OCR (SSO), password-protected wallet in the OCR (PKS), and Oracle Key Vault as the key store (OKV).

Examples

The following are examples of the use of `acfsutil encr info`.

Example 16-74 Using the `acfsutil encr info` command

```
# /sbin/acfsutil encr info -m /acfsmounts/acfs1

# /sbin/acfsutil encr info -m /acfsmounts/acfs1
                        -r /acfsmounts/acfs1/myfiles
```

16.10.2 `acfsutil encr init`

Purpose

Creates storage for encryption keys.

Syntax and Description

```
acfsutil encr init -h
acfsutil encr init [-p ]
```

`acfsutil encr init -h` displays help text and exits.

[Table 16-80](#) (page 16-119) contains the options available with the `acfsutil encr init` command.

Table 16-80 Options for the `acfsutil encr init` command

Option	Description
<code>-p</code>	Creates PKCS (password-protected) storage for keys.

The `acfsutil encr init` command must be run before any other encryption `acfsutil` commands can be run. This command must be run once for each cluster on which Oracle ACFS encryption is run.

If you plan to use Oracle Key Vault as the key store, then you still must run `acfsutil encr init` first.

See Also:

Oracle Key Vault Administrator's Guide for information about Oracle Key Vault

If the `-p` option is specified, you must provide a password when prompted. The password must conform to the format that is described in "[acfsutil sec init](#)" (page 16-93)".

If the `-p` option is not specified, a single sign-on (SSO) wallet is created.

Only a user with root or system administrator privileges can run this command.

Examples

The following is an example of the use of `acfsutil encr init`.

Example 16-75 Using the `acfsutil encr init` command

```
# /sbin/acfsutil encr init
```

16.10.3 `acfsutil encr off`

Purpose

Disables encryption for an Oracle ACFS file system, directories, or individual files.

Syntax and Description

```
acfsutil encr off -h
acfsutil encr off -m mount_point [[-r] path [ path ...]]
```

`acfsutil encr off -h` displays help text and exits.

[Table 16-82](#) (page 16-121) contains the options available with the `acfsutil encr off` command.

Table 16-81 Options for the `acfsutil encr off` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-r</code>	Specifies to disable encryption recursively under an existing directory identified by <i>path</i> .
<i>path</i>	Specifies the absolute or relative path of a directory. Multiple path values are allowed.

This command cannot be run on security realm-protected files.

Only an administrator can run this command on an Oracle ACFS file system (`-m` option without a *path* specified). When the `-m` option is specified without a *path*, all the files under the mount point are decrypted.

The *path* option can specify a path to a file or directory in a read-write snapshot. If the `-r` option is specified with the command on the root directory, the command does not transverse the snapshots under the `.ACFS` directory. If a decryption operation is specified at the file system level, then the operation does not process files and directories of snapshots in the `.ACFS/snaps/` directory.

Only a user with root or system administrator privileges can run this command to disable encryption on a file system. The file owner can also run this command to disable encryption on a directory or file.

Examples

The following are examples of the use of `acfsutil encr off`.

Example 16-76 Using the `acfsutil encr off` command

```
# /sbin/acfsutil encr off -m /acfsmounts/acfs1
```



```
# /sbin/acfsutil encr off -m /acfsmounts/acfs1
-r /acfsmounts/acfs1/myfiles
```

16.10.4 acfsutil encr on

Purpose

Encrypts an Oracle ACFS file system, directories, or individual files.

Syntax and Description

```
acfsutil encr on -h
acfsutil encr on -m mount_point
[-a {AES} -k {128|192|256}] [[-r] path [path...]]
```

`acfsutil encr on -h` displays help text and exits.

[Table 16-82](#) (page 16-121) contains the options available with the `acfsutil encr on` command.

Table 16-82 Options for the `acfsutil encr on` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-a <i>algorithm</i></code>	Specifies the encryption algorithm type for a directory or file. Advanced Encryption Standard (AES) is the only encryption algorithm supported for this release.
<code>-k <i>key_length</i></code>	Specifies the encryption key length for a directory or file.
<code>-r</code>	Specifies encryption recursively under existing directory folder identified by <i>path</i> .
<i>path</i>	Specifies the absolute or relative path of a directory. Multiple path values are allowed.

This command cannot be run on realm-protected files.

The default values for the `-a` and `-k` are determined by the volume parameters specified when `acfsutil encr set` was run. To set the key length at the volume level, use the `acfsutil encr set` command.

Only an administrator can run this command on an Oracle ACFS file system (`-m` option without a *path* specified). When the `-m` option is specified without a *path*, all the files under the mount point are encrypted.

The *path* option can specify a path to a file or directory in a read-write snapshot. If the `-r` option is specified with the command on the root directory, the command does not transverse the snapshots under the `.ACFS` directory. If an encryption operation is specified at the file system level, then the operation does not process files and directories of snapshots in the `.ACFS/snaps/` directory.

When you run `acfsutil encr on` with the `-r` option, the command encrypts the specified directory recursively, but does not enable encryption on the file system level.

Only a user with root or system administrator privileges can run this command to enable encryption on a file system. The file owner can also run this command to enable encryption on a directory or file.

Examples

The following are examples of the use of `acfsutil encr on`.

Example 16-77 Using the `acfsutil encr on` command

```
# /sbin/acfsutil encr on -m /acfsmounts/acfs1

# /sbin/acfsutil encr on -m /acfsmounts/acfs1
-a AES -k 128 -r /acfsmounts/acfs1/myfiles
```

16.10.5 acfsutil encr rekey

Purpose

Generates a new key and re-encrypts volume or file.

Syntax and Description

```
acfsutil encr rekey -h
acfsutil encr rekey -m mount_point
    {-f [-r] path [path...] |-v } [-a {AES} -k {128|192 |256}]
```

`acfsutil encr rekey -h` displays help text and exits.

[Table 16-83](#) (page 16-122) contains the options available with the `acfsutil encr rekey` command.

Table 16-83 Options for the `acfsutil encr rekey` command

Option	Description
<code>-m <i>mount_point</i></code>	Specifies the directory where the file system is mounted.
<code>-f [-r] <i>path</i> ...</code>	Generates a new file encryption key for the specified path and then encrypts the data with the new key. If <code>-r</code> is specified, the rekey operation is performed recursively under <i>path</i> . <i>path</i> specifies the absolute or relative path of a directory. Multiple path values are allowed.
<code>-v</code>	Generates a new volume encryption key (VEK) for the specified mount point and then encrypts all the file encryption keys in file system with the new key. Prompts for the wallet password because the wallet must be accessed to store the new VEK. The generated key is stored in the key store that was previously configured with the <code>acfsutil encr init</code> command.
<code>-a <i>algorithm</i></code>	Specifies the algorithm. Advanced Encryption Standard (AES) is the only encryption supported for this release.
<code>-k <i>key_length</i></code>	Specifies the key length for the directory or file specified by <i>path</i> .

This command cannot be run on security realm-protected files.

The default values for the `-a` and `-k` are determined by the volume parameters specified when `acfsutil encr set` was run.

The *path* option can specify a path to a file or directory in a read-write snapshot. If the *-r* option is specified with the command on the root directory, the command does not transverse the snapshots under the *.ACFS* directory. If a rekey operation is specified at the file system level, then the operation does not process files and directories of snapshots in the *.ACFS/snaps/* directory.

If Oracle Key Vault is the key store for the file system, then the Oracle Key Vault home environmental variable (*OKV_HOME*) must be set when using the *-v* option to generate a new volume key. If the client was configured to use a password with Oracle Key Vault, then the same password must be entered when prompted.

See Also:

Oracle Key Vault Administrator's Guide for information about Oracle Key Vault

Only a user with root or system administrator privileges can run this command with the *-v* option. The file owner can also run this command with the *-f* option to rekey encryption on the directory or file.

Examples

The following are examples of the use of `acfsutil encr rekey`.

Example 16-78 Using the `acfsutil encr rekey` command

```
# /sbin/acfsutil encr rekey -m /acfsmounts/acfs1 -v

# /sbin/acfsutil encr rekey -m /acfsmounts/acfs1 -f
                        -r /acfsmounts/acfs1/myfiles
```

16.10.6 acfsutil encr set

Purpose

Sets or changes encryption parameters for an Oracle ACFS file system.

Syntax and Description

```
acfsutil encr set -h
acfsutil encr set [ [-a {AES} -k {128|192|256}] [-e] | -u ] -m mount_point
```

`acfsutil encr set -h` displays help text and exits.

[Table 16-84](#) (page 16-123) contains the options available with the `acfsutil encr set` command.

Table 16-84 Options for the `acfsutil encr set` command

Option	Description
<code>-a algorithm</code>	Specifies the algorithm. Advanced Encryption Standard (AES) is the default value and the only encryption supported for this release. The algorithm must be specified if <code>-k</code> is specified.
<code>-k {128 192 256}</code>	Specifies the key length. The key length is set at the volume level. The default is 192. Must be specified if <code>-a</code> is specified.

Table 16-84 (Cont.) Options for the `acfsutil encr set` command

Option	Description
<code>-e</code>	Specifies to use Oracle Key Vault as the key store.
<code>-u</code>	Backs out encryption. Decrypts all encrypted files in the file system and reverts the file system to the state before <code>acfsutil encr set</code> was run on the file system. If security is being used, then this command can only be run after security has been backed out. To remove security, refer to " acfsutil sec prepare (page 16-96)".
<code>-m mount_point</code>	Specifies the directory where the file system is mounted.

Before running the `acfsutil encr set` command, you must first run the `acfsutil encr init` command.

The `acfsutil encr set` command configures encryption parameters for a file system, transparently generates a volume encryption key, and stores that the generated key in the key store that was previously configured with the `acfsutil encr init` command.

If auditing is initialized on a cluster, this command also enables an Oracle ACFS encryption auditing source on the file system. The actions performed when enabling this audit source are the same as those done when the `acfsutil audit enable` command is run directly. For more information, refer to "[acfsutil audit enable](#) (page 16-82)".

In addition `acfsutil encr set` creates the `mount_point/.Security/encryption/logs/` directory that contains the log file (`encr-hostname_fsid.log`) that collects auditing and diagnostic data.

Password requirements when storing the key are dependent on how the encryption key storage was configured. If `-p` was specified with `acfsutil encr init`, then a password is required to run this command.

Before using the `-e` option to specify Oracle Key Vault as the key store, Oracle Key Vault must be configured first. If you want to choose Oracle Key Vault as the key store for the file system, then the Oracle Key Vault home environmental variable (`OKV_HOME`) must be set when running the command with the `-e` option. If the client was configured to use a password with Oracle Key Vault, then the same password must be entered when prompted.

See Also:

Oracle Key Vault Administrator's Guide for information about configuring Oracle Key Vault

The `acfsutil encr set -u` command is not allowed if any snapshots exist in the file system.

Only a user with root or system administrator privileges can run the `acfsutil encr set` command.

Examples

The following example shows the use of `acfsutil encr set` command.

Example 16-79 Using the acfsutil encr set command

```
# /sbin/acfsutil encr set -a AES -k 256 -m /acfsmounts/acfs1
```

16.10.7 acfsutil keystore migrate**Purpose**

Migrates the encryption keystore.

Syntax and Description

```
acfsutil keystore migrate -h
acfsutil keystore migrate [-p ]
```

`acfsutil keystore migrate -h` displays help text and exits.

[Table 16-80](#) (page 16-119) contains the options available with the `acfsutil keystore migrate` command.

Table 16-85 Options for the acfsutil keystore migrate command

Option	Description
-p	Converts the encryption keystore from a SSO wallet to a PKCS wallet.

If the `-p` option is specified, `acfsutil keystore migrate` converts the SSO wallet to a PKCS wallet. If the `-p` option is not specified, `acfsutil keystore migrate` converts the PKCS wallet to a SSO wallet.

Only a user with root or system administrator privileges can run this command.

Examples

The following is an example of the use of `acfsutil keystore migrate`.

Example 16-80 Using the acfsutil keystore migrate command

```
# /sbin/acfsutil keystore migrate
```

16.11 Oracle ACFS Command-Line Tools for Snapshots

This topic provides a summary of commands to manage Oracle ACFS snapshots.

[Table 16-86](#) (page 16-125) lists Oracle ACFS snapshot commands with brief descriptions. For an overview of Oracle ACFS snapshots, refer to [About Oracle ACFS Snapshots](#) (page 11-11).

For more information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Table 16-86 Summary of commands for Oracle ACFS snapshots

Command	Description
<code>acfsutil snap convert</code> (page 16-126)	Converts the type of an existing snapshot.

Table 16-86 (Cont.) Summary of commands for Oracle ACFS snapshots

Command	Description
acfsutil snap create (page 16-127)	Creates a snapshot of an Oracle ACFS file system or an existing snapshot.
acfsutil snap delete (page 16-130)	Deletes a snapshot of an Oracle ACFS file system.
acfsutil snap duplicate apply (page 16-131)	Updates a duplicate snapshot.
acfsutil snap duplicate create (page 16-132)	Creates a duplicate of an existing snapshot.
acfsutil snap info (page 16-134)	Displays information about Oracle ACFS file system snapshots.
acfsutil snap quota (page 16-136)	Sets the quota limit for a specified snapshot.
acfsutil snap remaster (page 16-138)	Remasters a file system with a specified snapshot.
acfsutil snap rename (page 16-139)	Renames a snapshot.

16.11.1 acfsutil snap convert

Purpose

Converts the type of an existing snapshot image from read-only to read-write or read-write to read-only.

Syntax and Description

```
acfsutil snap convert -h
acfsutil snap convert [-r|-w] snapshot mount_point
```

`acfsutil snap convert -h` displays help text and exits.

[Table 16-87](#) (page 16-126) contains the options available with the `acfsutil snap convert` command.

Table 16-87 Options for the acfsutil snap convert command

Option	Description
<code>-r</code>	Converts snapshot to read-only snapshot.
<code>-w</code>	Converts snapshot to read-write snapshot.
<i>snapshot</i>	Specifies the name of the snapshot to be converted.
<i>mount_point</i>	Specifies the directory where the file system is mounted.

`acfsutil snap convert` converts the type of an existing snapshot image from a read-only snapshot image to a read-write snapshot image, or from a read-write snapshot image to a read-only snapshot image.

The type option (`-r` or `-w`) is required for the convert operation. Specifying a type parameter that matches the type of the existing snapshot image results in no

conversion and no failure. An 11.2 read-only snapshot image can be converted to a read-write snapshot image, but this conversion results in an update of the Oracle ACFS on-disk structure version. After the Oracle ACFS on-disk structure version has been updated, the updated version is not compatible with previous Oracle ACFS 11.2 versions.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-81](#) (page 16-127) shows the use of the `acfsutil snap convert` command.

Example 16-81 Using the `acfsutil snap convert` command

```
$ acfsutil snap convert -w midday_test1 /acfsmounts/critical_apps
$ acfsutil snap convert -r midday_test2 /acfsmounts/critical_apps
```

16.11.2 `acfsutil snap create`

Purpose

Creates a read-only or read-write snapshot of an Oracle ACFS file system or an existing snapshot.

Syntax and Description

```
acfsutil snap create -h
acfsutil snap create [-r|-w] [-c] [-p parent_snapshot] snapshot mount_point
```

`acfsutil snap create -h` displays help text and exits.

[Table 16-88](#) (page 16-127) contains the options available with the `acfsutil snap create` command.

Table 16-88 Options for the `acfsutil snap create` command

Option	Description
<code>-r</code>	Creates a read-only snapshot. This is the default setting.
<code>-w</code>	Creates a read-write snapshot.
<code>-c</code>	Specifies that duplicate specific data within the source file system or snapshot should be copied to the new snapshot. This option can only be used when the source of the snapshot to be created is the result of a <code>snap duplicate apply</code> operation.
<i>snapshot</i>	Specifies a name for the snapshot. The name provided must be a valid directory name. The <code>.ACFS/snaps</code> directory itself cannot be snapped.
<code>-p <i>parent_snapshot</i></code>	Specifies the name of an existing snapshot image within the Oracle ACFS file system.
<i>mount_point</i>	Specifies the directory where the file system is mounted.

`acfsutil snap create` creates a read-only or read-write snapshot of the Oracle ACFS file system mounted on `mount_point` or creates a snapshot of an existing snapshot if the `-p` option is specified. You can specify `-r` for read-only or `-w` for read-write; read-only is the default if neither `-r` or `-w` are specified.

The read-write snapshot enables the fast creation of an Oracle ACFS snapshot image that can be both read and written without impacting the state of the Oracle ACFS file system hosting the snapshot images. You can use the read-write functionality for testing new versions of application software or running test scenarios on production file data without modifying the original file system.

The `-c` option enables you to make a copy of a file system or snapshot whose last update was performed by an `acfsutil snap duplicate apply` operation. This updated file system or snapshot is known as the *apply target*. The data written to an *apply target* comes from the snapshot duplication stream of an `acfsutil snap duplicate create` operation.

An *apply target* contains metadata that identifies the stream that updated the target. This metadata ensures that any future `acfsutil snap duplicate apply` operation for the *apply target* has a starting point that matches the endpoint of the previous stream that updated the target. The following list summarizes the effects of the `-c` option.

- If the `-c` option is not specified when an *apply target* is copied, then the snapshot duplication metadata is cleared, which prevents the new snapshot from being a target of future apply operations.
- If `-c` is specified, then metadata is preserved, which enables the new snapshot to be used as the target of future `acfsutil snap duplicate apply` operations, exactly as the source file system or snapshot could be used.

The `-c` option produces a read-write snapshot, and may not be used with the `-r` or `-w` options.

Snapshots are not separate file systems. The snapshot appears in the `.ACFS/snaps/snapshot` directory and initially is a complete replica of the file system at the time the snapshot command was given. A read-only snapshot continues to preserve that initial point-in-time view. A read-write snapshot can be modified by updates written directly to the files located within the `.ACFS/snaps/snapshot` hierarchy.

A symbol link (symlink) file is a pointer to a path specified at the time that the symlink is created. For any existing symlink in the source of the `acfsutil snap create` command, a new link is made in the snapshot; however, the stored path to the symlink target is not modified during snap creation time. When an absolute path is used as the target of the link, the link target of the snapshot still refers to the absolute path, outside the newly-created snapshot namespace. Future modifications to the targeted file can occur through the snapshot symlink, not to a new target inside the snapshot namespace. The contents of the targeted file can also be changed outside of the snapshot. Because of this behavior, you should use relative path targets for all symlinks that are within the same file system.

Any user can access the snapshot directory by specifying the path name. However, the `.ACFS` directory itself is hidden from directory listings of the root of the file system. This prevents recursive commands, such as `rm -rf` or `acfsutil tag set -r`, from the root of the file system inadvertently operating on snapshot files.

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 snap info` or `acfsutil info fs` command.

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 `acfsutil snap create` operation fails if the cluster is in rolling migration.

Creation from an existing snapshot is not permitted if there are:

- Any snapshots present in the file system that were created with the ADVM compatibility set to less than 12.1
- Any snapshots of the file system that were created after ADVM compatibility was set to 12.1 but while 11.2 snapshots existed
- Any snapshot deletion cleanup operations still running in the background for the above conditions

You can display pending snapshot operations with the `acfsutil snap info` command.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-82](#) (page 16-129) shows the use of the `acfsutil snap create` command to create snapshots on a Linux platform.

Example 16-82 Using the `acfsutil snap create` command on Linux

```
$ acfsutil snap create -w midday_test1 /acfsmounts/critical_apps
$ acfsutil snap create -w midday_test2 /acfsmounts/critical_apps
$ /sbin/acfsutil snap create payroll_report1 /acfsmounts/critical_apps
$ /sbin/acfsutil snap create payroll_report2 /acfsmounts/critical_apps
$ ls /acfsmounts/critical_apps/.ACFS/snaps
midday_test1 midday_test2 payroll_report1 payroll_report2
```

[Example 16-83](#) (page 16-129) shows the use of the `acfsutil snap create` command to create snapshots on a Windows platform. This example shows the creation of a snapshot from an existing snapshot.

Example 16-83 Using the `acfsutil snap create` command on Windows

```
C:\oracle> acfsutil snap create /w snap_1 e:
C:\oracle> acfsutil snap create /w snap_1-1 /p snap_1 e:
```

See Also:

- [About Oracle ACFS Snapshots](#) (page 11-11) for information about Oracle ACFS snapshots
- [Oracle ACFS Disk Space Usage](#) (page 17-2) for information about the number of snapshots supported
- [acfsutil snap info](#) (page 16-134) for information about displaying pending snapshots operations
- [acfsutil info fs](#) (page 16-149) for information about displaying the space used for snapshots
- [Oracle ACFS Command-Line Tools for Snapshots](#) (page 16-125) for information about the `acfsutil snap duplicate apply` and `acfsutil snap duplicate create` commands

16.11.3 acfsutil snap delete

Purpose

Deletes a snapshot of an Oracle ACFS file system.

Syntax and Description

```
acfsutil -h snap delete
acfsutil snap delete snapshot mount_point
```

`acfsutil -h snap delete` displays help text and exits.

[Table 16-89](#) (page 16-130) contains the options available with the `acfsutil snap delete` command.

Table 16-89 Options for the `acfsutil snap delete` command

Option	Description
<i>snapshot</i>	Specifies a name for the snapshot.
<i>mount_point</i>	Specifies the directory where the file system is mounted.

`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 succeeds even if there are open files or directory references; however, storage associated with the snapshot is not released until all such references are closed.

The disk space used by the snapshot being deleted is released by a background task after the completion of the `acfsutil snap delete` command or after the last close of a file in the snapshot. If one of these background threads is running to clean up a deleted snapshot, then the `acfsutil snap info` command shows a pending delete operation. For information about the `acfsutil snap info` command, refer to "[acfsutil snap info](#) (page 16-134)".

The `acfsutil snap delete` operation fails if the cluster is in rolling migration.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-84](#) (page 16-131) shows the use of the `acfsutil snap delete` command.

Example 16-84 Using the `acfsutil snap delete` command

```
$ /sbin/acfsutil snap delete midday_test1 /acfsmounts/critical_apps
acfsutil snap delete: Snapshot operation is complete.
```

16.11.4 `acfsutil snap duplicate apply`

Purpose

Applies a previously created duplicated snapshot to an Oracle ACFS file system.

Syntax and Description

```
acfsutil snap duplicate apply -h
acfsutil snap duplicate apply [-b] [-d trace_level] [snap_name] mount_point
```

`acfsutil snap duplicate apply -h` displays help text and exits.

[Table 16-90](#) (page 16-131) contains the options available with the `acfsutil snap duplicate apply` command.

Table 16-90 Options for the `acfsutil snap duplicate apply` command

Option	Description
<code>-b</code>	Specifies that a backup snapshot is generated after a successful completion of the command
<code>-d trace_level</code>	Specifies the trace level setting [0..6] for the command execution.
<code>snap_name</code>	Specifies that the named snapshot of that file system is the target. The specified snapshot must exist and must be writable. If the <code>snap_name</code> is not specified, the target for the apply operation is the file system currently mounted at the mount point
<code>mount_point</code>	Specifies the directory where the duplicate snapshot is applied.

The `acfsutil snap duplicate apply` command uses the snapshot duplication stream produced by `acfsutil snap duplicate create` to update the specified target file system or snapshot. The stream of data is read from the standard input.

If the target has never been applied to by a `acfsutil snap duplicate apply` operation, this is an initial apply, and the target must be initially empty. If the target has been applied to by a previous apply operation, this is an incremental apply, and the initial contents of the target must match the contents of the older snapshot used to produce the snapshot duplication stream.

If the contents of the target no longer match the older snapshot for the stream, the command returns an error and the stream is not applied. The target must not be updated by other means while the apply operation is in progress. If another update occurs, the command returns an error and stops applying the snapshot duplication

stream. If the data stream being applied is identical to the last data stream applied to the specified target, the command exits without writing anything to the target with a status value of 2.

If the `-b` option is specified, a backup snapshot is taken of the target after the apply operation completes the update. Because the snapshot taken for an apply operation is automatically deleted when the next apply operation completes, the snapshot present is the one from the last-completed apply operation. The backup snapshot provides a known and consistent version of the target for the last successful apply operation and a comparison point to identify the in-progress changes made by a `acfsutil snap duplicate apply` operation.

A backup snapshot has a name of the form `SDBACKUP_secCnt_sourceSnap` or `SDBACKUP_secCnt_sourceSnap_targetSnap`. In the name string, `secCnt` is a local timestamp (number of seconds since the Unix epoch) indicating the point at which the target was captured by the backup snapshot. `sourceSnap` indicates the snapshot from the source file system that was applied by the last completed `acfsutil snap duplicate apply` operation. `targetSnap`, if present, indicates the snapshot name in the target file system that was updated by the last completed `acfsutil snap duplicate apply` operation.

System administrator (`root`) privileges are required to use this command.

The `acfsutil snap duplicate apply` command produces an exit status of 0 on success, 1 if a fatal error occurs, or 2 if the data stream read matches the last previous data stream applied to the specified target.

For information about displaying snapshot information, refer to "[acfsutil snap info](#) (page 16-134)".

Examples

[Example 16-85](#) (page 16-132) shows the use of the `acfsutil snap duplicate apply` command. Because the command expects to read a binary data stream from its standard input, this command is normally used in a pipeline with the `acfsutil snap duplicate create` command, as shown in the example.

Example 16-85 Using the `acfsutil snap duplicate apply` command

```
$ /sbin/acfsutil snap duplicate create new_snapshot_name /acfsmounts/critical_apps |
\  
    /sbin/acfsutil snap duplicate apply /acfsmounts/ca_backups
```

16.11.5 `acfsutil snap duplicate create`

Purpose

Creates a duplicate snapshot of an existing snapshot.

Syntax and Description

```
acfsutil snap duplicate create -h  
acfsutil snap duplicate create [-r] [-d trace_level] [-p parent_snapshot]  
                                [-i old_snapshot] new_snapshot mount_point
```

`acfsutil snap duplicate create -h` displays help text and exits.

[Table 16-91](#) (page 16-133) contains the options available with the `acfsutil snap duplicate create` command.

Table 16-91 Options for the `acfsutil snap duplicate create` command

Option	Description
<code>-r</code>	Specifies to mark the snap duplication stream as a restart stream. This option is intended for use on the second or later attempt to create and apply a given stream. This option might be necessary if an earlier attempt terminated prematurely.
<code>-d trace_level</code>	Specifies the trace level setting [0..6] for the command execution.
<code>-p parent_snapshot</code>	Specifies the name of a parent snapshot image of which both the <code>old_snapshot</code> and <code>new_snapshot</code> are children. This option ensures that both snapshots are children (<i>snap-of-snap</i>) of the same snapshot. If the <code>-p</code> option is not specified, then either the <code>new_snapshot</code> must be a child of the <code>old_snapshot</code> , or the two snapshots must be snapshots of the base file system at the specified mount point.
<code>-i old_snapshot</code>	Specifies the name of an existing snapshot.
<code>new_snapshot</code>	Specifies a name for the new snapshot.
<code>mount_point</code>	Specifies the directory where the file system is mounted. This is the mount point of the file system of which the snapshots were taken.

The `acfsutil snap duplicate create` command externalizes snapshot-related data in a form suitable for use by `snap duplicate apply` to create a different instance of a compatible snapshot. This externalized data is called a snapshot duplication stream.

If only the new snapshot name is specified, then this command encodes the entire contents of the named snapshot (the difference between an empty snapshot and this snapshot). This is an initial `snap duplicate` operation.

If both the old snapshot name and new snapshot name are given, this command encodes only the differences between the two snapshots. This is an incremental `snap duplicate` operation.

In either case, this command writes the snapshot duplication stream to its standard output.

System administrator (`root`) privileges are required to use this command.

The `acfsutil snap duplicate create` command produces an exit status of 0 on success or 1 if a fatal error occurs.

For information about displaying snapshot information, refer to "[acfsutil snap info](#) (page 16-134)".

Examples

[Example 16-86](#) (page 16-134) shows the use of the `acfsutil snap duplicate create` command. Because the command writes a binary data stream to its standard output, this command is normally used in a pipeline with the `acfsutil snap duplicate apply` command, as shown in the example.

Example 16-86 Using the `acfsutil snap duplicate create` command

```
$ /sbin/acfsutil snap duplicate create new_snapshot_name /acfsmounts/critical_apps |
\  
    /sbin/acfsutil snap duplicate apply /acfsmounts/ca_backups
```

16.11.6 acfsutil snap info**Purpose**

Displays information about Oracle ACFS file system snapshots.

Syntax and Description

```
acfsutil snap info -h
acfsutil snap info [-t] [snapshot] mount_point
```

`acfsutil snap info -h` displays help text and exits.

[Table 16-92](#) (page 16-134) contains the options available with the `acfsutil snap info` command.

Table 16-92 Options for the `acfsutil snap info` command

Option	Description
<code>-t</code>	Displays a tree structure for the specified snapshot and mount point.
<code>snapshot</code>	Specifies a name for the snapshot.
<code>mount_point</code>	Specifies the directory where the file system is mounted.

The `acfsutil snap info` command displays information about an individual specified snapshot or all snapshots in the specified Oracle ACFS file system.

The snapshot name, snapshot type (RO or RW), parent name, creation date and time, and space used are displayed. The parent name is either the mount point, or the parent snapshot if the snapshot was created from an existing snapshot. Also included in this information are the flag values indicating that either a snapshot is replication-specific or is in use as the target of a `snap duplicate apply` operation. If the snapshot is replication-specific, then the flag value `REPL` appears. If the snapshot is in use by `snap duplicate apply`, then the flag value `DUP` appears. These flag values are displayed in the output of the command following the flag value indicating that the snapshot is read-only (RO) or read-write (RW).

The snapshot space usage amount includes snapshot metadata. If all the files are deleted from a snapshot, some metadata still remains and that amount is displayed with `acfsutil snap info`.

When using an accelerator volume, the reported snapshot space usage amount can include space that is in use on the accelerator volume.

To reclaim the disk space used by snapshots, delete the snapshots. Deleting the files within the snapshots may not reclaim this disk space.

The `acfsutil snap info` command shows the status of a snapshot that is in the process of being deleted as follows:

- When waiting on files to be closed:

```
snapshot_name (delete waiting for last close)
```

- When freeing the storage for the snapshot in the background:

```
snapshot_name (delete in progress)
```

Oracle ACFS caches and asynchronously updates snapshot storage usage statistics. If an Oracle ASM instance is killed, power is lost to the machine, or the Oracle ACFS file system goes offline for any reason, one of these cached updates may be lost. This loss can result in the display of a message similar to `Reported snapshot space usage is inaccurate` when `acfsutil snap info` is run.

Examples

[Example 16-87](#) (page 16-135) shows the use of the `acfsutil snap info` command on a Linux platform. The command output contains information about a read-only (RO) snapshot, a read-write (RW) snapshot with quota information, a snapshot in `delete waiting for last close` status, and a snapshot in `delete in progress` status. After a delete operation completes on a snapshot and the background process finishes the clean up of a deleted snapshot, messages about the delete operation do not appear in the output.

Example 16-87 Using the `acfsutil snap info` command on Linux

```
$ /sbin/acfsutil snap info /acfsmounts/critical_apps/
snapshot name:          midday_test (delete waiting for last close)
snapshot creation time: Tue Apr  5 10:52:55 2016

snapshot name:          payroll_report1
snapshot location:     /acfsmounts/critical_apps/.ACFS/snaps/payroll_report1
RO snapshot or RW snapshot: RW
parent name:           /acfsmounts/critical_apps/
snapshot creation time: Tue Apr  5 10:52:59 2016
storage added to snapshot: 95514624 ( 91.09 MB )

quota:                 104857600 ( 100.00 MB )
quota usage:           91 %

snapshot name:          midday_test2 (delete in progress)
snapshot creation time: Mon Apr 11 11:40:13 2016

snapshot name:          payroll_report2
snapshot location:     /acfsmounts/critical_apps/.ACFS/snaps/payroll_report2
RO snapshot or RW snapshot: RO
parent name:           midday_test2
snapshot creation time: Mon Apr 11 12:05:42 2016
storage added to snapshot: 1081344 ( 1.03 MB )

number of snapshots:  3 (active)
                    1 (delete in progress)
snapshot space usage: 102084608 ( 97.36 MB )

$ /sbin/acfsutil snap info -t /acfsmounts/critical_apps
/acfsmounts/critical_apps RW
midday_test                RW /acfsmounts/critical_apps
payroll_report1            RW /acfsmounts/critical_apps
```

```

midday_test2          RW   /acfsmounts/critical_apps
payroll_report2      RO   midday_test2

```

```

$ /sbin/acfsutil snap info -t midday_test2 /acfsmounts/critical_apps
midday_test2          RW
payroll_report2      RO   midday_test2

```

[Example 16-88](#) (page 16-136) shows the use of the `acfsutil snap info` command on a Windows platform.

Example 16-88 Using the `acfsutil snap info` command on Windows

```

C:\oracle> acfsutil snap info e:
snapshot name:          MIDDAY_TEST (delete waiting for last close)
snapshot creation time: Wed Apr 13 12:44:15 2016

snapshot name:          PAYROLL_REPORT1
snapshot location:     e:\.ACFS\snaps\PAYROLL_REPORT1
RO snapshot or RW snapshot: RW
parent name:           e:
snapshot creation time: Wed Apr 13 12:45:00 2016
storage added to snapshot: 18866176 ( 17.99 MB )

quota:                 104857600 ( 100.00 MB )
quota usage:           17 %

snapshot name:          MIDDAY_TEST2 (delete in progress)
snapshot creation time: Wed Apr 13 12:48:07 2016

snapshot name:          PAYROLL_REPORT2
snapshot location:     e:\.ACFS\snaps\PAYROLL_REPORT2
RO snapshot or RW snapshot: RO
parent name:           MIDDAY_TEST2
snapshot creation time: Wed Apr 13 12:50:34 2016
storage added to snapshot: 1081344 ( 1.03 MB )

number of snapshots:  3 (active)
                    1 (delete in progress)
snapshot space usage: 362557440 ( 345.76 MB )

```

```

C:\oracle> acfsutil snap info /t e:
e:
MIDDAY_TEST          RW   e:
PAYROLL_REPORT1     RW   e:
MIDDAY_TEST2        RW   e:
PAYROLL_REPORT2     RO   MIDDAY_TEST2

```

```

C:\oracle> acfsutil snap info /t MIDDAY_TEST2 e:
MIDDAY_TEST2        RW
PAYROLL_REPORT2     RO   MIDDAY_TEST2

```

16.11.7 acfsutil snap quota

Purpose

Sets the quota limit for a specified snapshot.

Syntax and Description

```
acfsutil -h snap quota
```

```
acfsutil snap quota [+|-]n[K|M|G|T|P] snapshot mount_point
```

`acfsutil -h snap quota` displays help text and exits.

The following table contains the options available with the `acfsutil snap quota` command.

Table 16-93 Options for the `acfsutil snap quota` command

Option	Description
[+ -]n K M G T P	Specifies the quota limit size for the Oracle ACFS snapshot 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. 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.
<i>snapshot</i>	Specifies a name of the snapshot.
<i>mount_point</i>	Specifies the directory where the file system is mounted.

The `acfsutil snap quota` command sets the quota limit for a specified snapshot on the specified Oracle ACFS file system.

The quota can be larger than the size of the file system because a file system can be resized or automatically resized.

The quota limit specified does not reserve space. The quota only establishes an upper limit on the amount of storage that can be added to the snapshot.

Only snapshots created after ADVN compatibility (`COMPATIBLE.ADVN`) has been set to 1.2.2 or greater have storage accounting tracking and the ability to have quotas set.

Snapshots can be converted from read-write (RW) to read-only (RO) and vice versa. Both RW and RO snapshots can have a quota.

There is a grace period and advanced warnings for quota usage. A warning message is written to both the alert log and the console if the requested storage allocation causes the local storage usage to be near or exceeding more than 90% of the quota. If an allocation exceeds the quota, then the write or other modifying operation fails as `EDQUOT`, or it may fail as `ENOSPC`, depending whether the quota limit is set larger than the size of the file system.

System administrator (`root`) privileges are required to use this command.

You can display snapshot information, including the quota limit and quota usage percentage, with the `acfsutil snap info` command. For information about `acfsutil snap info`, refer to "[acfsutil snap info](#) (page 16-134)".

Examples

The following example shows the use of the `acfsutil snap quota` command.

Example 16-89 Using the `acfsutil snap quota` command

```
$ /sbin/acfsutil snap quota +500M snap_test /acfsmounts/acfs1
```

16.11.8 acfsutil snap remaster

Purpose

Remasters an Oracle ACFS file system with a snapshot.

Syntax and Description

```
acfsutil snap remaster -h
acfsutil snap remaster { snapshot | -c } -f volume
```

`acfsutil snap remaster -h` displays help text and exits.

[Table 16-94](#) (page 16-138) contains the options available with the `acfsutil snap remaster` command.

Table 16-94 Options for the `acfsutil snap remaster` command

Option	Description
<code>-c</code>	Continues the interrupted remaster operation. After the command has started, the remastering must run to completion.
<code>-f</code>	Forces the operation and does not prompt for questions.
<i>snapshot</i>	Specifies the name of an existing snapshot.
<i>volume</i>	Specifies specifies the primary Oracle ADVM volume device path for the file system.

This command accepts a snapshot name or the `-c` option.

The specified snapshot is used for the base of the file system, replacing the current base of the file system. Using an existing read-only or read-write snapshot, the command reintegrates the snapshot extents into the base of the current file system creating a new remastered file system with the point in time state of the snapshot. The remastered file system is left with no snapshots. This operation must be done on an unmounted file system and it cannot be reversed.

Use the `-c` option, instead of the snapshot name, to complete an interrupted snapshot remastering.

You cannot remaster a file system on which replication, security, or encryption is running. You must terminate replication, security, or encryption before running the `acfsutil snap remaster` command. In addition, you cannot remaster a file system with plugins enabled.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-90](#) (page 16-139) shows the use of the `acfsutil snap remaster` command.

Example 16-90 Using the acfsutil snap remaster command

```
$ /sbin/acfsutil snap remaster snapshot volume
```

16.11.9 acfsutil snap rename**Purpose**

Renames an Oracle ACFS snapshot.

Syntax and Description

```
acfsutil snap rename -h
acfsutil snap rename old_snapshot new_snapshot mount_point
```

`acfsutil snap rename -h` displays help text and exits.

[Table 16-95](#) (page 16-139) contains the options available with the `acfsutil snap duplicate create` command.

Table 16-95 Options for the acfsutil snap rename command

Option	Description
<i>old_snapshot</i>	Specifies the name of an existing snapshot.
<i>new_snapshot</i>	Specifies a name for the new snapshot.
<i>mount_point</i>	Specifies the directory where the file system is mounted.

The `acfsutil snap rename` command is used to safely rename an existing snapshot in the filesystem.

While a snapshot with open files can be renamed, applications accessing files in that snapshot experience the same behavior as operating in a non-snap directory tree in which one of the parent directories was renamed.

Renaming snapshots used in conjunction with pluggable databases (PDBs) is not supported at this time. Such a snapshot should not be renamed.

Replication-specific snapshots cannot be renamed as this may cause complications and inconsistencies with replication.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

For information about displaying snapshot information, refer to "[acfsutil snap info](#)" (page 16-134)".

Examples

[Example 16-91](#) (page 16-140) shows the use of the `acfsutil snap rename` command.

Example 16-91 Using the `acfsutil snap rename` command

```
$ /sbin/acfsutil snap rename old_snapshot_name new_snapshot_name /acfsmounts/
critical_apps
```

16.12 Oracle ACFS Command-Line Tools for Compression

This topic provides a summary of commands to manage Oracle ACFS compression.

[Table 16-96](#) (page 16-140) lists Oracle ACFS compression commands with brief descriptions. For an overview of Oracle ACFS compression, refer to [Oracle ACFS Compression](#) (page 11-28).

For more information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Table 16-96 Summary of commands for Oracle ACFS compression

Command	Description
acfsutil compress copy (page 16-140)	Copies and compresses a file or files in a directory.
acfsutil compress info (page 16-141)	Displays compression information about a file.
acfsutil compress off (page 16-142)	Enables Oracle ACFS file compression.
acfsutil compress on (page 16-142)	Disables Oracle ACFS file compression.

16.12.1 `acfsutil compress copy`

Purpose

Compresses and copies a file.

Syntax and Description

```
acfsutil compress copy -h
acfsutil compress copy [-v] [-f] {-c size source | [-n] {-r source | source [...]}}
                        target_dir
```

The following table contains the options available with the `acfsutil compress copy` command.

Table 16-97 Options for the `acfsutil compress copy` command

Option	Description
<code>-v</code>	Specifies verbose mode.
<code>-f</code>	Specifies to overwrite any existing copies.
<code>-c size source</code>	Copies a single source file and uses compression unit (CU) <i>size</i> for the copy. The format of <i>size</i> is <i>nK</i> . The CU sizes supported are: 4K, 8K, 16K, 32K, 64K and 128K bytes.
<code>-n source</code>	Specifies that copies created are not compressed.

Table 16-97 (Cont.) Options for the `acfsutil compress copy` command

Option	Description
<code>-r source</code>	Specifies to recursively copy files and directories starting from a single source directory.
<code>source [...]</code>	Specifies source files to be copy.
<code>target_dir</code>	Specifies the target directory to copy into.

Only files recognized as Oracle Database files that are normally compressed are compressed during the copy operation. With these files, the copy is compressed using the CU size stored in the Oracle ACFS metadata if it exists, and otherwise, the command uses the DB blocksize stored in the file.

If the file is recognized as a database file and the DB blocksize stored in the file is not valid as a CU size and the Oracle ACFS metadata is not available, then the copy is not compressed.

If the file is not recognized as a database file, the copy is not compressed.

The command does not copy over an existing files unless the `-f` option is used. If `-f` is not used, the copy operation terminates when the first existing target is found.

The command does not copy to an Oracle ACFS file system which is capable of supporting compression if compression is disabled unless the `-n` option is specified.

Example 16-92 Using the `acfsutil compress copy` command

```
$ /sbin/acfsutil compress copy -v -f my_file1 my_file2 /my_target_directory
```

16.12.2 `acfsutil compress info`

Purpose

Displays information about compressed files.

Syntax and Description

```
acfsutil compress info -h
acfsutil compress info pathname
```

The following table contains the options available with the `acfsutil compress info` command.

Table 16-98 Options for the `acfsutil compress info` command

Option	Description
<code>pathname</code>	Specifies the pathname to a file.

Example 16-93 Using the `acfsutil compress info` command

```
$ /sbin/acfsutil compress info /my_mount_point/testfile1
```

```
Compression Unit size: 8192
Disk storage used: ( 189.89 MB )
```

Disk storage saved: (810.20 MB)
Storage used is 18% of what the uncompressed file would use

16.12.3 acfsutil compress off

Purpose

Disables Oracle ACFS compression.

Syntax and Description

```
acfsutil compress off -h
acfsutil compress off mount_point
```

The following table contains the options available with the `acfsutil compress off` command.

Table 16-99 Options for the `acfsutil compress off` command

Option	Description
<i>mount_point</i>	Specifies the directory where the file system is mounted.

When running `acfsutil compress off`, the following apply:

- Files that are created after compression is disabled are created uncompressed.
- Any compressed files in the file system remain compressed.

Example 16-94 Using the `acfsutil compress off` command

```
$ /sbin/acfsutil compress off /my_mount_point
```

16.12.4 acfsutil compress on

Purpose

Enables Oracle ACFS compression.

Syntax and Description

```
acfsutil compress on -h
acfsutil compress on [-a algorithm] mount_point
```

The following table contains the options available with the `acfsutil compress on` command.

Table 16-100 Options for the `acfsutil compress on` command

Option	Description
<code>-a <i>algorithm</i></code>	Specifies the default compression algorithm. Currently only <code>lzo</code> is available and that is the default algorithm used.
<i>mount_point</i>	Specifies the directory where the file system is mounted.

When running `acfsutil compress on`, the following apply:

- Files created after compression is enabled are compressed by default.
- Any uncompressed files in the file system remain uncompressed.

Example 16-95 Using the `acfsutil compress on` command

```
$ /sbin/acfsutil compress on /my_mount_point
```

16.13 Oracle ACFS Command-Line Utilities

This topic provides a summary of the Oracle ACFS command-line utilities.

[Table 16-101](#) (page 16-143) lists the Oracle ACFS command-line utilities with brief descriptions.

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Table 16-101 Summary of Oracle ACFS command-line utilities

Command	Description
acfsutil accel replace (page 16-143)	Replaces an existing accelerator volume with a new one.
acfsutil compat get (page 16-144)	Displays the compatibility level of an Oracle ACFS file system.
acfsutil compat set (page 16-145)	Changes the compatibility level of an Oracle ACFS file system.
acfsutil defrag dir (page 16-146)	Defragments all the files in a specified directory.
acfsutil defrag file (page 16-146)	Defragments the specified files.
acfsutil info file (page 16-147)	Displays information for a file in an Oracle ACFS file system.
acfsutil info fs (page 16-149)	Displays detailed Oracle ACFS file system information.
acfsutil info id (page 16-154)	Displays Oracle ACFS file system information for a specified file identifier and mount point.
acfsutil plugin disable (page 16-155)	Disables the Oracle ACFS plug-in infrastructure.
acfsutil plugin enable (page 16-156)	Enables the Oracle ACFS plug-in infrastructure.
acfsutil plugin info (page 16-159)	Displays information about the Oracle ACFS plug-in infrastructure.
acfsutil registry (page 16-160)	Adds, deletes, or displays entries in the Oracle ACFS mount registry.
acfsutil rmfs (page 16-162)	Removes an Oracle ACFS file system.
acfsutil scrub (page 16-163)	Checks for and reports any inconsistencies in the metadata or file data.
acfsutil size (page 16-164)	Resizes an Oracle ACFS file system.
acfsutil tune (page 16-167)	Modifies or displays Oracle ACFS tunable parameters.
advmutil canonical (page 16-168)	Displays the canonical name of an Oracle ADVM volume.
advmutil tune (page 16-169)	Modifies or displays Oracle ADVM parameters.
advmutil volinfo (page 16-170)	Displays information about Oracle ADVM volumes.

16.13.1 `acfsutil accel replace`

Purpose

Replaces an existing accelerator volume with a new one.

Syntax and Description

```
acfsutil accel replace -h
acfsutil accel replace -a new_accel_volume primary_volume
```

`acfsutil accel replace -h` displays help text and exits.

The following table contains the options available with the `acfsutil accel replace` command.

Table 16-102 Options for the `acfsutil accel replace` command

Option	Description
<code>-a <i>new_accel_volume</i></code>	Specifies the name of the new accelerator volume.
<code><i>primary_volume</i></code>	Specifies the name of the primary volume with which the Oracle ACFS file system is mounted.

This command allows an administrator to replace an existing accelerator volume with a new accelerator volume. The command is useful if the current accelerator volume is full and cannot be grown, or to migrate an accelerator to a faster volume.

The file system must be unmounted on all nodes to use this command. The new accelerator must be at least as large as the existing accelerator. If new accelerator volume is larger than the existing accelerator volume, then the administrator should run `acfsutil resize -d new_accel_volume` after the file system is remounted following the `acfsutil accel replace` operation to enable Oracle ACFS to utilize the entire space of the new accelerator volume.

You must be the `root` user or an Oracle ASM administrator user to run this command.

The following example shows the use of the `acfsutil accel replace` command.

Example 16-96 Using the `acfsutil accel replace` command

```
$ /sbin/acfsutil accel replace -a new_accelerator_volume my_primary_volume
```

16.13.2 acfsutil compat get**Purpose**

Displays the Oracle ACFS compatibility level for the file system and the `COMPATIBLE.ADVM` level for the disk group.

Syntax and Description

```
acfsutil compat get -h
acfsutil compat get mount_point
```

`acfsutil compat get -h` displays help text and exits.

The following table contains the options available with the `acfsutil compat get` command.

Table 16-103 Options for the `acfsutil compat get` command

Option	Description
<code>mount_point</code>	Specifies the directory where the file system is mounted.

The `acfsutil compat get` command displays the Oracle ACFS compatibility level for the file system that may have been set by the `mkfs` command or the `acfsutil compat set` command. `acfsutil compat get` also displays the `COMPATIBLE.ADVM` setting for the disk group, which is default compatibility level for the file system if not otherwise explicitly set. For information about Oracle ASM disk group compatibility attributes, see "[Disk Group Compatibility](#) (page 4-50)".

The following example shows the use of the `acfsutil compat get` command.

Example 16-97 Using the `acfsutil compat get` command

```
$ /sbin/acfsutil compat get /acfsmounts/acfs1
```

16.13.3 `acfsutil compat set`

Purpose

Change the compatibility level of an Oracle ACFS file system.

Syntax and Description

```
acfsutil compat set -h
acfsutil compat set -c release_version mount_point
```

`acfsutil compat set -h` displays help text and exits.

The following table contains the options available with the `acfsutil compat set` command.

Table 16-104 Options for the `acfsutil compat set` command

Option	Description
<code>-c <i>release_version</i></code>	Specifies the release version for the Oracle ACFS file system determined by <code>mount_point</code> .
<code>mount_point</code>	Specifies the directory where the file system is mounted.

The value specified with the `-c release_version` option must be greater than or equal to the value of `COMPATIBLE.ADVM` for the disk group, and must be less than or equal to the running Oracle Grid infrastructure release version. If `-c release_version` is not specified, the value of `COMPATIBLE.ADVM` is used. After the compatibility is set, it cannot be downgraded. The `-c release_version` option can be used in situations where it is not possible, or desirable, to update `COMPATIBLE.ADVM` and `COMPATIBLE.ASM` for the disk group, but you want to use an Oracle ACFS feature that requires a compatibility increase. After the Oracle ACFS compatibility is updated, it is no longer possible to mount the file system with an older Oracle Grid infrastructure release. For information about Oracle ASM disk group compatibility attributes, see "[Disk Group Compatibility](#) (page 4-50)".

The following example shows the use of the `acfsutil compat set` command.

Example 16-98 Using the `acfsutil compat set` command

```
$ /sbin/acfsutil compat set -c 12.2 /acfsmounts/acfs1
```

16.13.4 acfsutil defrag dir**Purpose**

Defragments all the files in a specified directory.

Syntax and Description

```
acfsutil defrag dir -h
acfsutil defrag dir [-r] dir_path
```

`acfsutil defrag dir -h` displays help text and exits.

[Table 16-105](#) (page 16-146) contains the options available with the `acfsutil defrag dir` command.

Table 16-105 Options for the `acfsutil defrag dir` command

Option	Description
<code>-r</code>	Specifies recursive defragment operation.
<code>dir_path</code>	Specifies the directory where the files to be defragmented are located.

The `acfsutil defrag dir` command is used to defragment files in the specified subdirectory. The `-r` option is used to recursively defragment files in the subdirectories of the specified directory.

You must have write access to the files to be defragmented.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-99](#) (page 16-146) shows the use of the `acfsutil defrag dir` command.

Example 16-99 Using the `acfsutil defrag dir` command

```
$ /sbin/acfsutil defrag dir /acfsmounts/critical_apps/01jan1015
```

16.13.5 acfsutil defrag file**Purpose**

Defragments the specified files.

Syntax and Description

```
acfsutil defrag file -h
acfsutil defrag file [-i] file_path [file_path ... ]
```

`acfsutil defrag file -h` displays help text and exits.

[Table 16-106](#) (page 16-147) contains the options available with the `acfsutil defrag file` command.

Table 16-106 Options for the `acfsutil defrag file` command

Option	Description
<code>-r</code>	Estimates the reduction in number of extents possible. Requires read access to the file.
<code>file_path</code>	Specifies the file to be defragmented.

The `acfsutil defrag file` command defragments files. The `-i` option displays the estimated reduction in the number of extents for the files to be defragmented.

You must have write access to the files to be defragmented.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-100](#) (page 16-147) shows the use of the `acfsutil defrag file` command.

Example 16-100 Using the `acfsutil defrag file` command

```
$ /sbin/acfsutil defrag file /acfsmounts/critical_apps/testfile
```

16.13.6 `acfsutil info file`

Purpose

Displays information for a file in an Oracle ACFS file system.

Syntax and Description

```
acfsutil info file -h
acfsutil info file [ [-d] [-o acfs_extent_offset] [-u] | [-c] ] path [path ... ]
```

`acfsutil info file -h` displays help text and exits.

[Table 16-107](#) (page 16-147) contains the options available with the `acfsutil info file` command.

Table 16-107 Options for the `acfsutil info file` command

Option	Description
<code>-d</code>	Displays detailed extent information, from Oracle ACFS to the Oracle ASM devices in the disk group.
<code>-o acfs_extent_offset</code>	<code>acfs_extent_offset</code> — Filters the extent display to the Oracle ACFS file offset specified.
<code>-c</code>	Display extent counts only.
<code>-u</code>	Display extent usage information.
<code>path</code>	Specifies the full path name of a file.

The `acfsutil info file` command displays information about a file in an Oracle ACFS file system. The information includes the extent map which details the locations of the blocks comprising the file and the used versus allocated storage for a file. Tagging information is also displayed for a file.

The `-u` option adds an additional column to the standard `acfsutil info file` output indicating whether or not the extent is inherited between a snapshot parent and child. The `-c` option displays the total extent count and the number of primary and accelerator volume extent nodes, while omitting the details of each individual extent. The output from the `-c` option is useful when you want to display the number of extents, and when using an accelerator, to ensure that the metadata is situated on the accelerator volume.

You must have read access to the specified file to run `acfsutil info file`.

Examples

The following is an example of `acfsutil info file`.

Example 16-101 Using the `acfsutil info file` command

```
$ /sbin/acfsutil info file /usmmnt/.ACFS/snaps/s1/afile
/usmmnt/.ACFS/snaps/s1/afile
  flags:      File
  inode:     18014398509482029
  owner:     root
  group:     root
  size:      13013 ( 12.71 KB )
  allocated: 24576 ( 24.00 KB )
  hardlinks: 1
  device index: 1
  major, minor: 202,64
  access time: Thu Jan  7 14:13:24 2016
  modify time: Thu Jan  7 14:13:24 2016
  change time: Thu Jan  7 14:13:24 2016
  extents:
    -offset ----length | -dev -----offset
           0      24576 |   1      142188544
  extent count: 1
  granularity level: -1

$ /sbin/acfsutil info file -u /usmmnt/.ACFS/snaps/s1/afile
/usmmnt/.ACFS/snaps/s1/afile
  flags:      File
  inode:     18014398509482029
  owner:     root
  group:     root
  size:      13013 ( 12.71 KB )
  allocated: 24576 ( 24.00 KB )
  hardlinks: 1
  device index: 1
  major, minor: 202,64
  access time: Thu Jan  7 14:13:24 2016
  modify time: Thu Jan  7 14:13:24 2016
  change time: Thu Jan  7 14:13:24 2016
  extents:
    extent count: 1
    primary volume extent nodes: 0
    accelerator volume extent nodes: 0
    granularity level: -1
```

```

$ /sbin/acfsutil info file -c /usmmnt/.ACFS/snaps/s1/afile
/usmmnt/.ACFS/snaps/s1/afile
  flags:      File
  inode:      18014398509482029
  owner:      root
  group:      root
  size:       13013 ( 12.71 KB )
  allocated:  24576 ( 24.00 KB )
  hardlinks:  1
  device index: 1
  major, minor: 202,64
  access time: Thu Jan  7 14:13:24 2016
  modify time: Thu Jan  7 14:13:24 2016
  change time: Thu Jan  7 14:13:24 2016
  extents:
    -offset ----length | -dev -----offset | inherited
          0      24576 |    1      142188544 | Yes
  extent count: 1
  added extent count: 0
  inherited extent count: 1
  granularity level: -1

```

16.13.7 acfsutil info fs

Purpose

Displays detailed Oracle ACFS file system information.

Syntax and Description

```

acfsutil info fs -h
acfsutil info fs [-1] [{-o option_item | -s [interval[count]][-d]] | -f [-v]]
                    [mount_point]

```

`acfsutil info fs -h` displays help text and exits.

[Table 16-108](#) (page 16-149) contains the options available with the `acfsutil info fs` command.

Table 16-108 Options for the `acfsutil info fs` command

Option	Description
-1	Displays symbolic links for an Oracle ACFS file system specified by the mount point or displays symbolic links for all mounted Oracle ACFS file systems if a mount point is not specified.

Table 16-108 (Cont.) Options for the `acfsutil info fs` command

Option	Description
-o	<p data-bbox="690 321 1375 373">Displays the specific file system option item from the following list:</p> <ul data-bbox="690 388 1375 1419" style="list-style-type: none"> <li data-bbox="690 388 1375 499">• <code>allvolumes [mount_point primary_volume_device auxiliary_volume_device]</code> — Lists all volumes that are associated with a file system. <li data-bbox="690 506 1375 533">• <code>autoresize</code> — 1 if automatic resize is enabled, 0 if not. <li data-bbox="690 539 1375 592">• <code>autoresizeincrement</code> — Lists the automatic resize increment for the file system. <li data-bbox="690 598 1375 625">• <code>autoresizemax</code> — Lists the automatic resize maximum. <li data-bbox="690 632 1375 743">• <code>auxvolume {mount_point primary_volume_device}</code> — Lists the auxiliary volumes associated with the file system. Can be run against a mount point or device. <li data-bbox="690 749 1375 802">• <code>available</code> - Length of time the file system has been available. <li data-bbox="690 808 1375 835">• <code>compatacfs</code> — Lists the ACFS compatibility setting <li data-bbox="690 842 1375 869">• <code>compatadv</code> — Lists the ADV <li data-bbox="690 875 1375 928">• <code>diskgroup</code> - Name of the disk group associated with <code>path</code>. <li data-bbox="690 934 1375 987">• <code>fetasize</code> - File entry table size which is the amount of storage consumed by the on disk metadata for inodes. <li data-bbox="690 993 1375 1045">• <code>freespace</code> - Amount of free space available in the file system. <li data-bbox="690 1052 1375 1079">• <code>isadv</code> - 1 if an Oracle ADV <li data-bbox="690 1085 1375 1169">• <code>isauxvolume {volume_device}</code> — 1 if a volume is an auxiliary volume or 0 if not. Must be run against a volume device. <li data-bbox="690 1176 1375 1228">• <code>isavailable</code> - 1 if the file system is available, 0 if file system is not available. <li data-bbox="690 1234 1375 1287">• <code>iscompression</code> — State of compression on the file system. 1 for enabled, 0 for disabled. <li data-bbox="690 1293 1375 1346">• <code>iscorrupt</code> - 1 if the file system is corrupt, 0 if file system is not corrupt. <li data-bbox="690 1352 1375 1419">• <code>isencryption</code> — State of encryption on the file system, 0 for disabled, 1 for enabled.

Table 16-108 (Cont.) Options for the `acfsutil info fs` command

Option	Description
<code>-o</code> (continued)	<ul style="list-style-type: none"> <code>ismountpoint</code> - 1 if <i>path</i> is the mount point, 0 if <i>path</i> is not a mount point. <code>isreplprimary</code> - 1 if primary file system, 0 if not. <code>isreplstandby</code> - 1 if standby file system, 0 if not. <code>issecurity</code> — State of security on the file system, 0 for disabled, 1 for enabled. <code>metablocksize</code> — Size of metadata blocks on the file system. <code>mountpoints</code> - Mount points associated with file systems. <code>primaryvolume</code> - Path name of the device associated with file system. <code>redundancy</code> - Type of redundancy used by the Oracle ADVM volume. <code>replication</code> - 1 if file system has been initialized for replication as the primary or standby file system, 0 if not initialized for replication. <code>resizeincrement</code> - Allocation unit size in bytes of the Oracle ADVM volume used by <i>path</i>. <code>stripewidth</code> - Stripe width for the Oracle ADVM volume used by <i>path</i>. <code>stripecolumns</code> - Stripe set columns for the Oracle ADVM volume used by <i>path</i>. <code>totalspace</code> - Size of the file system. <code>volumes</code> - Number of volumes associated with the file system.
<code>-s</code> [<i>interval</i>] [<i>count</i>] [<i>-d</i>]	<p>Displays file system statistics.</p> <p>The optional <i>interval</i> parameter specifies the amount of time in seconds between each report. The first report contains statistics for the time since the file system mount. Each subsequent report contains statistics collected during the interval since the previous report.</p> <p>The optional <i>count</i> parameter can be specified with the <i>interval</i> parameter. If the <i>count</i> parameter is specified, the value of <i>count</i> determines the number of reports generated at interval seconds apart.</p> <p>If the <i>interval</i> parameter is specified without the <i>count</i> parameter, the command generates reports continuously until you break out of the command with <code>Ctrl+c</code>.</p> <p>The optional <code>-d</code> option displays more detailed statistical information.</p>
<code>-f</code> [<i>-v</i>]	Displays file system fragmentation. The <code>-v</code> option displays additional fragmentation information.
<i>mount_point</i>	Specifies the directory where the file system is mounted. Not a valid option for replication data.

`acfsutil info fs` displays information and statistics about Oracle ACFS file systems. If compression, encryption, security, or replication has been initiated on a file system, additional flags and status information are displayed.

The `-o option_item` option displays file system information for the `option_item` specified.

The `-s` option displays the amount and rate of change currently on a file system for the node that the command is run on.

With no option specified, the command displays file system information that 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 Oracle ACFS snapshots, the space consumed by snapshots in the file system, and the optional name or volume label associated with the file system. The possible file system states are displayed in the `flags` line. These 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 the file system should be checked and repaired at the earliest possible convenience to correct a detected inconsistency. For example, run the `fsck` command on Linux or the `acfschkdsk` command on Windows to check and repair the file system. Ensure that you run the command in repair mode to correct the problem. For more information, refer to "[fsck](#) (page 16-26)" and "[acfschkdsk](#) (page 16-50)".

When viewing space usage values with `acfsutil info fs` on Windows systems, the values may differ from sizes in Windows folders. The mechanism used by **Folder Properties** on Windows systems only accounts for files and should be considered an approximate value.

Any user can run `acfsutil info fs`.

Examples

The following are examples of the use of `acfsutil info fs`.

[Example 16-102](#) (page 16-152) displays information about an Oracle ACFS file system in a Linux environment.

Example 16-102 Using the `acfsutil info fs` command on Linux

```
$ /sbin/acfsutil info fs
/primary
  ACFS Version: 11.2.0.2.0
  flags:       MountPoint,Available,Replication
  mount time:  Mon Oct 25 12:11:03 2010
  volumes:     1
  total size:  5368709120
  total free:  4144230400
  primary volume: /dev/asm/pvol-74
    label:
    flags:       Primary,Available,ADVM
    on-disk version: 40.0
    allocation unit: 4096
    major, minor: 252, 37889
    size:        5368709120
    free:        4144230400
    ADVM diskgroup REPLDG
```



```

        ADVM resize increment: 268435456
        ADVM redundancy:      unprotected
        ADVM stripe columns:  4
        ADVM stripe width:    131072
    number of snapshots: 0
    snapshot space usage: 0
    replication status: primary

/standby
    ACFS Version: 11.2.0.2.0
    flags:        MountPoint,Available,Replication
    mount time:   Mon Oct 25 12:11:03 2010
    volumes:     1
    total size:   5368709120
    total free:   5263945728
    primary volume: /dev/asm/svol-74
        label:
        flags:        Primary,Available,ADVM
        on-disk version: 40.0
        allocation unit: 4096
        major, minor:  252, 37890
        size:          5368709120
        free:          5263945728
        ADVM diskgroup REPLDG
        ADVM resize increment: 268435456
        ADVM redundancy:      unprotected
        ADVM stripe columns:  4
        ADVM stripe width:    131072
    number of snapshots: 0
    snapshot space usage: 0
    replication status: standby

$ /sbin/acfsutil info fs -o mountpoints,replication
/primary
1
/standby
1

$ /sbin/acfsutil info fs -o mountpoints,isreplprimary
/primary
1
/standby
0

$ /sbin/acfsutil info fs -o mountpoints,isreplstandby
/primary
0
/standby
1

```

[Example 16-103](#) (page 16-153) illustrates the use of `acfsutil info fs` in a Windows environment.

Example 16-103 Using the `acfsutil info fs` command on Windows

```

C:\oracle>acfsutil info fs /o freespace c:\oracle\acfsmounts\acfs1
968667136

C:\oracle>acfsutil info fs /o totalspace c:\oracle\acfsmounts\acfs1
1073741824

C:\oracle>acfsutil info fs /o volumes c:\oracle\acfsmounts\acfs1

```

1

```
C:\oracle>acfsutil info fs /o primaryvolume c:\oracle\acfsmounts\acfs1
asm-volume1-311
```

```
C:\oracle>acfsutil info fs /o diskgroup c:\oracle\acfsmounts\acfs1
DATA
```

```
C:\oracle>acfsutil info fs /o redundancy c:\oracle\acfsmounts\acfs1
mirror
```

Example 16-104 (page 16-154) illustrates the use of `acfsutil info fs` with `-s` option to display the current amount and rate of change on a file system.

Example 16-104 Using the `acfsutil info fs` command with the `-s` option

```
$ /sbin/acfsutil info fs -s /acfsmounts/acfs1
amount of change since mount:      359.22 MB
average rate of change since mount: 3 KB

$ /sbin/acfsutil info fs -s 60 4 /acfsmounts/acfs1
amount of change since mount:      359.22 MB
average rate of change since mount: 3 KB/s

amount of change: 15.02 MB   rate of change: 256 KB/s
amount of change:  9.46 MB   rate of change: 161 KB/s
amount of change:  7.32 MB   rate of change: 125 KB/s
amount of change:  6.89 MB   rate of change: 117 KB/s
...
```

16.13.8 acfsutil info id

Purpose

Displays Oracle ACFS file system information for a specified file identifier number and mount point.

Syntax and Description

```
acfsutil info id -h
acfsutil info id file_identifier mount_point
```

`acfsutil info id -h` displays help text and exits.

Table 16-109 (page 16-154) contains the options available with the `acfsutil info id` command.

Table 16-109 Options for the `acfsutil info id` command

Option	Description
<i>mount_point</i>	Specifies the directory where the file system is mounted.
<i>file_identifier</i>	Specifies the Oracle ACFS file identifier number reported by the Oracle ACFS driver. The file identifier number should be specified in decimal format.

You can use `acfsutil info id` 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 is an example of `acfsutil info id`.

Example 16-105 Using the `acfsutil info id` command

```
$ /sbin/acfsutil info id 117 /acfsmounts/acfs1
```

16.13.9 acfsutil plugin disable

Purpose

Disables the Oracle ACFS plug-in infrastructure for an Oracle ACFS file system.

Syntax and Description

```
acfsutil plugin disable -h
acfsutil plugin disable mount_point
```

`acfsutil plugin disable -h` displays help text and exits.

[Table 16-110](#) (page 16-155) contains the options available with the `acfsutil plugin disable` command.

Table 16-110 Options for the `acfsutil plugin disable` command

Option	Description
<code>mount_point</code>	Specifies the directory where the file system is mounted.

`acfsutil plugin disable` disables the collection of Oracle ACFS plug-in metrics and the associated communication with a plug-in application.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-106](#) (page 16-155) shows the use of the `acfsutil plugin disable` command. After monitoring has completed for the Oracle ACFS file metric data for the file system mounted on the `/humanresources`, the command in [Example 16-106](#) (page 16-155) disables the subsequent collection of summary data in the Oracle ACFS driver.

Example 16-106 Disabling the Oracle ACFS driver

```
# /sbin/acfsutil plugin disable /humanresource
```

See Also:

- [About Using Oracle ACFS Command-Line Tools](#) (page 16-2) for information about running Oracle ACFS `acfsutil` commands
- [Oracle ACFS Plugins](#) (page 11-43) for an overview of Oracle ACFS plug-ins
- [Oracle ACFS Plug-in Generic Application Programming Interface](#) (page 17-8) for information about the Oracle ACFS plug-in application programming interface

16.13.10 acfsutil plugin enable

Purpose

Enables the Oracle ACFS plug-in infrastructure for an Oracle ACFS file system.

Syntax and Description

```
acfsutil plugin enable -h
acfsutil plugin enable -m metrictype [-t tag, ...]
                        [-i interval[s|m]] mount_point
```

`acfsutil plugin enable -h` displays help text and exits.

[Table 16-111](#) (page 16-156) contains the options available with the `acfsutil plugin enable` command.

Table 16-111 Options for the `acfsutil plugin enable` command

Option	Description
<code>-t <i>tag, ...</i></code>	An optional list of tag names selecting Oracle ACFS files that you want to be monitored. The maximum number of tags is 12. The names are separated by commas. The default is to monitor all files in the Oracle ACFS file system.
<code>-m <i>metrictype</i></code>	Selects the message payload to be transmitted from Oracle ACFS to the plug-in module with each call-out message. Oracle ACFS provides the <code>acfsmetric1</code> pre-defined metric type.

Table 16-111 (Cont.) Options for the `acfsutil plugin enable` command

Option	Description
<code>-i interval [s m]</code>	<p>Selects posting message delivery and the posting interval. The Oracle ACFS plug-in mechanism posts a metrics message during each specified message interval.</p> <p>An interval is the maximum amount of time which can pass before metrics are posted. The application blocks in the API call to retrieve metrics until a posting occurs.</p> <p>The interval can be specified as an integer value in <i>s</i> seconds or <i>m</i> minutes. The default is minutes. The maximum value allowed for an interval is 60 minutes or 3600 seconds.</p> <p>If an interval is specified, then Oracle ACFS is requested to post metrics messages to the plug-in application on an interval basis.</p> <p>If an interval is not specified, then the application polls for metrics updates from the Oracle ACFS file system. This is the default action.</p>
<code>mount_point</code>	Specifies the directory where the file system is mounted.

`acfsutil plugin enable` enables an Oracle ACFS file system for application plug-in service. Oracle ACFS plug-ins are enabled for single nodes only, either a standalone host or one or more individual nodes of a Grid Cluster. After enabled for plug-ins, the Oracle ACFS file system begins collecting metric information for either the set of specified tagged files or all files in the Oracle ACFS file system, and transmits callout messages to the plug-in application based upon the selected message payload and delivery parameters. Metrics are reset after every delivery to the application (whether through polling or posting).

In general, metrics are collected from the point of view of a user application. I/O for metadata or internal Oracle ACFS activities; such as snapshot copy-on-writing, encrypting, and replicating files; are not included. Metrics are collected for user I/O to both files in the original file system and snapshot files. Memory mapping of files that results in I/O operations is recorded in the metrics. Replication related files under the `.ACFS/repl` directory are not included in the metrics.

Oracle ACFS provides the `acfsmetric1` pre-defined metric type.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-107](#) (page 16-157) shows the use of the `acfsutil plugin enable` command to provide additional storage usage metrics to a monitoring application.

Example 16-107 Enabling the Oracle ACFS driver for storage visibility: poll model

```
# /sbin/acfsutil plugin enable -m acfsmetric1 -t HRDATA /humanresource
```

An Oracle ACFS file system mounted on `/humanresources` is to be enabled for plug-in service and configured with the objective of providing additional storage usage metrics to a monitoring application. The message payload type is `acfsmetric1`, the files to be monitored are human resource files that have been tagged with the Oracle ACFS file tag `HRDATA`, and the message delivery type is `poll`

which means that the application plug-in polls to request metric data update messages. No interval is specified with polled message delivery.

When the command in [Example 16-107](#) (page 16-157) completes, the Oracle ACFS file system mounted on the `/humanresources` mount point is enabled for plug-in communication and begins collecting Oracle ACFS file access metrics for files tagged with `HRDATA`. Oracle ACFS maintains metrics collecting a summary of read and write activity. Each time the associated application plug-in module polls for the metrics, Oracle ACFS sends a message with the data as defined in the `ACFS_METRIC1` structure that is a summary of the selected Oracle ACFS activity since either the Oracle ACFS driver plug-in functionality was enabled or since the last delivery of metric data.

[Example 16-108](#) (page 16-158) shows the use of the `acfsutil plugin enable` command to monitor movie file updates.

Example 16-108 Enabling the Oracle ACFS driver for storage visibility: post on interval model

```
# /sbin/acfsutil plugin enable -m acfsmetric1 -t FILECONTENT -i 120s /moviemods
```

An Oracle ACFS file system mounted on `/moviemods` is enabled for plug-in service and is configured to record and deliver storage usage messages to a monitoring application. On an interval basis, the Oracle ACFS file system sends a wakeup on a driver event on which the application API call is waiting. The application then collects a summary of read and write activity on the files.

The message payload type is `acfsmetric1` and the files to be monitored are movie files that have been tagged with the `FILECONTENT` tag. The message delivery type is `post`, which means that Oracle ACFS posts messages containing the metrics to the plug-in application for the specified interval of 120 seconds.

When the command in [Example 16-108](#) (page 16-158) completes, the Oracle ACFS file system mounted on the `/moviemods` mount point is enabled for plug-in communication and begins collecting Oracle ACFS storage visibility metrics for files tagged with `FILECONTENT`. On each interval, Oracle ACFS posts one or more messages to the application plug-in containing the storage visibility information.

See Also:

- [About Using Oracle ACFS Command-Line Tools](#) (page 16-2) for information about running Oracle ACFS `acfsutil` commands
 - [Oracle ACFS Plugins](#) (page 11-43) for an overview of Oracle ACFS plug-ins
 - [Oracle ACFS Pre-defined Metric Types](#) (page 17-8) for information about the pre-defined metric types
 - [Oracle ACFS Plug-in Generic Application Programming Interface](#) (page 17-8) for information about the Oracle ACFS plug-in application programming interface
-

16.13.11 acfsutil plugin info

Purpose

Displays information about the Oracle ACFS plug-in infrastructure for an Oracle ACFS file system.

Syntax and Description

```
acfsutil plugin info -h
acfsutil plugin info mount_point
```

`acfsutil plugin info -h` displays help text and exits.

[Table 16-112](#) (page 16-159) contains the options available with the `acfsutil plugin info` command.

Table 16-112 Options for the `acfsutil plugin info` command

Option	Description
<i>mount_point</i>	Specifies the directory where the file system is mounted.

`acfsutil plugin info` displays information about the state of the Oracle ACFS plug-in mechanism.

Administrator privileges are required to use this command or you must be a member of the Oracle ASM administrator group.

Examples

[Example 16-109](#) (page 16-159) shows the use of the `acfsutil plugin info` command.

Example 16-109 Displaying information about the Oracle ACFS driver

```
# /sbin/acfsutil plugin info /humanresource
```

```
Plug-in status: ENABLED
Metric type: acfsmetric1
Enabled tags: HRDATA
Delivery method: Poll
Post interval (seconds):
```

```
# /sbin/acfsutil plugin info /moviemods
```

```
Plug-in status: ENABLED
Metric type: acfsmetric1
Enabled tags: FILECONTENT
Delivery method: Post
Post interval (seconds): 120
```

See Also:

- [About Using Oracle ACFS Command-Line Tools](#) (page 16-2) for information about running Oracle ACFS `acfsutil` commands
- [Oracle ACFS Plugins](#) (page 11-43) for an overview of Oracle ACFS plug-ins
- [Oracle ACFS Plug-in Generic Application Programming Interface](#) (page 17-8) for information about the Oracle ACFS plug-in application programming interface

16.13.12 acfsutil registry**Purpose**

Adds, deletes, or displays entries in the Oracle ACFS mount registry.

Syntax and Description

```
acfsutil registry -h
acfsutil registry
acfsutil registry -a [-f] [-o moptions] [-n { nodes|all } ]
    [-u user] [-t description] device mount_point
acfsutil registry -c {device | mount_point} [-o moptions]
    [-n { nodes|all } ] [-u user] [-t description]
acfsutil registry -d {device | mount_point}
acfsutil registry -l [device | mount_point]
acfsutil registry -m deviceacfsutil registry -r
```

`acfsutil registry -h` displays help text and exits.

If no options are specified, the command displays all the Oracle ACFS file systems that are configured for automatic start.

[Table 16-113](#) (page 16-160) contains the options available with the `acfsutil registry` command.

Table 16-113 Options for the `acfsutil registry` command

Option	Description
-a	<p>Add the device, mount point, and associated <i>moptions</i> to the Oracle ACFS mount registry. The Oracle ADVM volume device specified must exist on the local node to add the information to the mount registry.</p> <p>The arguments represent all the information needed to mount the file system. At Oracle ACFS startup time these file systems are automatically mounted.</p> <p>Duplicate device entries are not allowed. Duplicate mount points are allowed but must include the <code>-n</code> option for disjoint node-specific mounts.</p>
-c	<p>Changes an existing registry entry for the specified device or mount point.</p>

Table 16-113 (Cont.) Options for the `acfsutil registry` command

Option	Description
<code>-d</code>	Deletes the specified device or mount point from the Oracle ACFS mount registry. If a mount point is specified and it is not unique in the Oracle ACFS mount registry, the command fails and you must run the command again specifying the device.
<code>-f</code>	This option 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.
<code>-l [device mount_point]</code>	Without specifying a device or mount point, the <code>-l</code> option lists all the file systems currently in the registry on a single line, with fields separated by a colon (:). With a device specified, lists information about the device in the Oracle ACFS registry. With a mount point specified, lists information about the mount point in the Oracle ACFS registry.
<code>-m device</code>	Lists the registered mount point, if one exists, associated with the specified device. The mount point is only returned if the Oracle ACFS file system has been registered or has been previously mounted.
<code>-n {nodes all}</code>	This option, used in combination with <code>-a</code> and <code>-c</code> , 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 for use when mounting the file system. Valid for Linux, Solaris, and AIX. Used in combination with <code>-a</code> and <code>-c</code> . For specific <code>-o options</code> , refer to the <code>mount</code> command for each operating system. All the options for each operating system are available except the <code>all</code> option.
<code>-r</code>	Displays all registered file systems, not just file systems with <code>auto_start=always</code> .
<code>-t description</code>	Adds a description of the mount. Used in combination with <code>-a</code> and <code>-c</code> . For example: <code>-t "HR Exports share"</code> . This description can be seen later when looking at configured resources.
<code>-u user</code>	Specifies a user that is allowed to mount or unmount (start and stop) the file system. This option, used in combination with <code>-a</code> and <code>-c</code> , is useful for creating a registered file system that can be started or stopped by someone other than root.
<code>mount_point</code>	Specifies the directory where the file system is mounted.
<code>device</code>	Specifies an Oracle ACFS device file that has been formatted.

`acfsutil registry` adds, deletes, or displays 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](#) (page 3-28)".

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.

Note:

Oracle ACFS registration (`acfsutil registry`) is not supported in an Oracle Restart (standalone) configuration, which is a single-instance (non-clustered) environment.

See "[About the Oracle ACFS Mount Registry](#) (page 11-11)".

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 16-110 Using the `acfsutil registry` command

```
$ /sbin/acfsutil registry -a /dev/asm/volume1-123 /acfsmounts/acfs1
$ /sbin/acfsutil registry -m /dev/asm/volume1-123
$ /sbin/acfsutil registry -d /dev/asm/volume1-123
```

16.13.13 `acfsutil rmfs`

Purpose

Removes an Oracle ACFS file system.

Syntax and Description

```
acfsutil rmfs -h
acfsutil rmfs device
acfsutil rmfs device -b
```

`acfsutil rmfs -h` displays help text and exits.

[Table 16-114](#) (page 16-162) contains the options available with the `acfsutil rmfs` command.

Table 16-114 Options for the `acfsutil rmfs` command

Option	Description
<code>device</code>	Specifies an Oracle ACFS device file that has been formatted.
<code>-b</code>	Specifies batch mode. No user interaction is required.

You can use `acfsutil rmfs` to remove an Oracle ACFS that is dismounted. When the command is run, 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](#) (page 3-28)".

After `acfsutil rmfs` runs 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 16-111 Using the `acfsutil rmfs` command

```
$ /sbin/acfsutil rmfs /dev/asm/volume1-123
```

16.13.14 `acfsutil scrub`

Purpose

Checks for and reports any inconsistencies in the metadata or file data.

Syntax and Description

```
acfsutil scrub [-h] [-n] [-m] [-q] [-e number] [-b ext1[,...]] [-z] [-a] [-p level]
               [-t level] {file | directory}
```

`acfsutil -h scrub` displays help text and exits.

The following table contains the options available with the `acfsutil scrub` command.

Table 16-115 Options for the `acfsutil scrub` command

Option	Description
<code>-n</code>	Specifies not to scrub the directory recursively.
<code>-m</code>	Specifies to scrub the metadata only.
<code>-q</code>	Specifies quiet mode. Only report inconsistencies to STDOUT.
<code>-e <i>number</i></code>	Specifies the number of inconsistencies encountered before program terminates.
<code>-b <i>ext1,ext2, ...</i></code>	Specifies not to scrub files with the specified extensions.
<code>-z</code>	Specifies not to scrub the Oracle ACFS snapshot directory.
<code>-a</code>	Specifies to scrub extensions skipped by default (<code>.dbf</code> , <code>.arc</code>).
<code>-p <i>level</i></code>	Specifies I/O load level, from 0 (lowest) to 6 (highest).
<code>-t <i>level</i></code>	Specifies trace level verbosity, from 0 (terse) to 6 (verbose).
<i>file</i>	Specifies the full path name of a file.

Table 16-115 (Cont.) Options for the `acfsutil scrub` command

Option	Description
<code>directory</code>	Specifies the full path name of a directory.

On Linux and AIX systems, the `acfsutil scrub` command checks the mirror consistency of Oracle ACFS directory metadata and file user data and metadata.

For paths with metadata inconsistencies, the scrubber displays the inconsistent path.

For file paths with user data inconsistencies, the scrubber provides the following information:

- Pathname
- Byte offset of inconsistency
- Byte length of inconsistency
- ADVM volume byte offset of inconsistency

You must have `root` privileges to run `acfsutil scrub`.

Examples

The following example illustrates the use of the `acfsutil scrub` command. The first command scrubs only the metadata of the specified file. The second command scrubs the specified directory and all nested files, except for snapshots, `txt` files, and `pdf` files. The third command scrubs the specified directory and all nested files with the power level set to 5 and trace level set to 3.

Example 16-112 Using the Oracle ACFS `acfsutil scrub` command

```
$ /sbin/acfsutil scrub -m /acfsmnt/textdoc1.txt

$ /sbin/acfsutil scrub -b txt,pdf -z /acfsmnt/dir1

$ /sbin/acfsutil scrub -p 5 -t 3 /acfsmnt/dir1
```

16.13.15 `acfsutil size`

Purpose

Resizes an Oracle ACFS file system.

Syntax and Description

```
acfsutil size -h
acfsutil size [ [+|-]n[K|M|G|T|P] ] [ -a n[K|M|G|T] ] [ -x n[K|M|G|T] ] [-d
volume_device] mount_point
```

`acfsutil size -h` displays help text and exits.

[Table 16-116](#) (page 16-165) contains the options available with the `acfsutil size` command.

Table 16-116 Options for the `acfsutil size` command

Option	Description
<code>[+ -]n</code> <code>K M G T P</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. 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.
<code>-an</code>	Specifies the increment by which an automatic resize increases the file system when the amount of free space falls below the amount specified by the increment. A nonnegative numeric value must be provided, and may optionally be followed by K (Kilobytes), M (Megabytes), G (Gigabytes), or T (Terabytes) to specify a value in kilobytes, megabytes, gigabytes, or terabytes. If the unit is omitted, the default is bytes. A value of 0 disables automatic resize. The minimum allowed increment is 1G.
<code>-xn</code>	Specifies the maximum size to which a file system automatically resizes. A nonnegative numeric value must be provided, and may optionally be followed by K (Kilobytes), M (Megabytes), G (Gigabytes), or T (Terabytes) to specify a value in kilobytes, megabytes, gigabytes, or terabytes. If the unit is omitted, the default is bytes. If this value is less than the current size of the file system then no error is returned, but automatic resize does not take place until the file system is shrunk (by manual action of the system administrator) below this value. A value of 0 indicates that no maximum size is configured, so the file system may continue to automatic resize until the underlying volume is out of space.
<code>-d volume_device</code>	Resizes only the specified device.
<code>mount_point</code>	Specifies the directory where the file system is mounted

At least one of the + or -, -a, or -x options must be specified.

The size value is rounded up based on the block size of the file system and the allocation unit of the Oracle ADVM volume device file. To determine the Oracle ADVM volume device resize increment, examine the `RESIZE_UNIT_MB` field in the `V $ASM_VOLUME` view, `Resize Unit` in the output of `asmcmd volinfo`, or `ADVM resize increment` in the output of `acfsutil info fs`.

There is no limit to the number of times a file system can be expanded for a disk group with the ADVM compatibility attribute set to `11.2.0.4` or higher.

For a disk group with the ADVM compatibility attribute set to less than `11.2.0.4`, there is a limit of 5 extents for the file system's internal storage bitmap. This causes any attempts to increase the file system to fail after it has been increased four or more times. However, if after increasing the file system four times or more times the file system size is decreased, then you may be able to increase the file system size again if the size of the increase is less than the size of the decrease. 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.

Note:

Oracle ACFS supports the shrinking of a file system, but just for those regions that only contain free blocks.

`root` or users who are members of the `asmadmin` group can run this command. For information about operating system group privileges, see "[About Privileges for Oracle ASM](#) (page 3-28)".

Manual Resize Option

The `+` or `-` option 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 ADVM volume file to match the new length that is specified. However, 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.

Automatic Resize Option

The automatic resize option (`-a` option) specifies an increment by which an Oracle ACFS file system grows automatically if the amount of available free space in the file system falls below the amount specified by the increment. You can also specify the maximum size (`-x` option) allowed when using the automatic resize option. The output of the `acfsutil info fs` command displays the automatic resize increment and maximum amounts. If only the `-x` option is provided without the `-a` option (and no automatic resize increment has previously been specified) then a warning is issued if the increment is still zero. If the increment had previously been set by an earlier invocation of the command, then the command succeeds. Automatic resize requires ADVM compatibility attribute set to 12.2 or higher.

If an automatic resize fails on a node for any reason, automatic resizing on the node is temporarily disabled. An error message is written to the USM event log and the OKS log. Automatic resizing may still take place on another node; however, if the failure is due to the underlying volume being out of free space then all nodes quickly disable automatic resize. The file system does continue to periodically try to automatic resize, at increasingly longer intervals. After a resize succeeds, such as when more space is added to the underlying Oracle ASM disk group, then automatic resize becomes re-enabled on the node. Note that other nodes may remain in the error state even after one node successfully auto-resizes the file system. Automatic resize can be manually re-enabled by either re-mounting the file system or running the `acfsutil size` command again with the `-a` or `-x` option to update automatic resize settings. When a file system has grown to the maximum size specified with `acfsutil size -x`, automatic resize remains enabled but does not grow the file system any further until the maximum is increased.

The `-d` argument is not compatible with the `-a` and `-x` automatic resize arguments. Unless a volume device is specified with `-d`, it is assumed that the caller is specifying an amount of storage that applies to the primary volume, which represents the size of the file system. When growing a file system with an accelerator volume, the

`acfsutil size` command determines if the accelerator size needs to be increased. If the accelerator must be larger and cannot be increased, the command fails and the primary volume remains at the same size. When shrinking a file system, only the primary volume is shrunk. If a volume device is specified with the `-d` option, then the size specified applies to either the primary volume or the accelerator volume, whichever is specified, and only that volume device is changed. The command warns the user if the resulting accelerator volume size is not large enough to accommodate the primary volume size.

When using an accelerator volume and automatic resize, the increment specified with the automatic resize `-a` argument applies to the primary volume; however, both the primary volume and the accelerator volume are monitored. Automatic resize grows the primary volume by the specified increment, or greater if needed. Automatic resize monitors and grows the accelerator volume as needed as well by smaller increments. When automatic resize needs to grow the primary volume, it first attempts to grow the accelerator if insufficient space remains on the accelerator to accommodate the new growth, and if that fails the primary volume remains at the same size.

In addition to being triggered because the free space falls below the configured automatic resize increment, an automatic resize also takes place under other conditions, such as:

- If the free space in the file system falls below 10% of the current file system size or 5GB, whichever is smaller.
- If a copy-on-write cannot find enough contiguous space in the file system to allocate new storage for the write.
- If the defragger cannot find enough contiguous space to allocate in order to defragment 8M worth of extents in a file.

These conditions protect against a large file system having enough total free space so that an automatic resize appears unnecessary, but the space is so fragmented that applications cannot use it.

When a non-sparse file is resized so that it grows by more than the configured auto-resize interval, the file system automatic resizes by a multiple of the automatic resize increment so that the resize request can be satisfied.

Examples

The following example shows the use of `acfsutil size`. This example increases the primary device file of `/acfsmounts/acfs1` file system by 500 MB.

Example 16-113 Using the `acfsutil size` command

```
$ /sbin/acfsutil size +500M /acfsmounts/acfs1
```

16.13.16 `acfsutil tune`

Purpose

The `acfsutil tune` command displays the value of a specific tunable parameter or all Oracle ACFS tunable parameters, or sets the value of a tunable parameter in a persistent manner on a particular node.

Syntax and Description

```
acfsutil tune -h
acfsutil tune [tunable_name]
acfsutil tune tunable_name=value
```

`acfsutil tune -h` displays help text and exits.

[Table 16-117](#) (page 16-168) contains the options available with the `acfsutil tune` command.

Table 16-117 Options for the `acfsutil tune` command

Option	Description
<code>tunable_name</code>	Specifies the name of the tunable parameter.
<code>value</code>	Specifies the value for a tunable parameter.

The Oracle ACFS tunable parameter `AcfsMaxOpenFiles` limits the number of open Oracle ACFS files on Windows and AIX. Normally you do not have to change the value of this tunable parameter; however, you may want to consider increasing the value if you have a large working set of files in your Oracle ACFS file systems.

The Oracle ACFS tunable parameter `AcfsMaxCachedFiles` sets the maximum number of closed files that remain cached in memory on Windows and AIX. Normally you do not have to change value of this tunable parameter; however, you many to consider changing the value to get better performance.

Changing a tunable 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 tunable parameter.

Examples

The first command prints Oracle ACFS tunable parameters. The second command changes the value of a tunable parameter.

Example 16-114 Using the `acfsutil tune` command

```
$ /sbin/acfsutil tune
$ acfsutil tune AcfsMaxOpenFiles=50000
```

16.13.17 advmutil canonical**Purpose**

`advmutil canonical` displays the canonical name of the specified Oracle ADVM device name.

Syntax and Description

```
advmutil -h
advmutil canonical volume_device
```

`advmutil -h` displays help text and exits.

[Table 16-118](#) (page 16-169) contains the options available with the `advmutil canonical` command.

Table 16-118 Options for the `advmutil canonical` command

Option	Description
<code>volume_device</code>	Specifies a string identifying an Oracle ADVM volume device.

There are several different formats that can identify an Oracle ADVM volume device, but a normalized, unambiguous (canonical) name should be used when the volume device name is used with other commands such as `SRVCTL`.

For example, on the Windows operating system you can use the following prefixes with a volume device name: `\\. \, \??\, \?\`

The `advmutil canonical` command would return the canonical name that another utility would recognize without having to strip off extra characters. The command would most likely be used in a script.

Examples

The following examples show the use of `advmutil canonical` on a Windows operating system. For each example, the command returns the canonical name of the volume device.

Example 16-115 Using `advmutil canonical`

```
[C:\]advmutil canonical asm-volume1-274
asm-volume1-274
```

```
[C:\]advmutil canonical \\. \asm-volume1-274
asm-volume1-274
```

```
[C:\]advmutil canonical \?\asm-volume1-274
asm-volume1-274
```

```
[C:\]advmutil canonical \??\asm-volume1-274
asm-volume1-274
```

16.13.18 `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]
```

`advmutil -h` displays help text and exits.

[Table 16-119](#) (page 16-170) contains the options available with the `advmutil tune` command.

Table 16-119 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 parameter that can be specified with `advmutil tune` is the maximum time in minutes for the deadlock timer (`deadlock_timer`).

You must be a privileged user to set a parameter.

Note:

The `deadlock_timer` parameter should only be set by Oracle Support Services.

Examples

The first example changes the maximum time in minutes for the deadlock timer. The second example queries the current setting of a parameter.

Example 16-116 Using `advmutil tune`

```
$ /sbin/advmutil tune deadlock_timer = 20

$ /sbin/advmutil tune deadlock_timer
deadlock_timer = 20 (0x14)
```

16.13.19 `advmutil volinfo`**Purpose**

`advmutil volinfo` displays information about Oracle ADVM volume devices.

Syntax and Description

```
advmutil -h
advmutil volinfo [-l][-L] [volume_device]
```

`advmutil -h` displays help text and exits.

[Table 16-120](#) (page 16-170) contains the options available with the `advmutil volinfo` command.

Table 16-120 Options for the `advmutil volinfo` command

Option	Description
<i>volume_device</i>	Specifies an optional volume device name.

Table 16-120 (Cont.) Options for the `advmutil volinfo` command

Option	Description
-l	Separates the Oracle ADVM volume device information by field descriptions and colons on one line.
-L	Separates the Oracle ADVM volume device information by spaces on one line.

`advmutil volinfo` displays information about Oracle ADVM volume devices in a list format by default. The `-l` option on Linux (`/l` on Windows) formats the display into a colon-separated string with field descriptions. The `-L` option on Linux (`/L` on Windows) formats the display into a space-separated string in a format suitable for scripting.

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 using the `-l` option. The third example displays information about the volume device using the `-L` option.

Example 16-117 Using `advmutil volinfo`

```
$ /sbin/advmutil volinfo /dev/asm/volume1-123
Device : /dev/asm/volume1-228
Interface Version: 1
Size (MB): 256
Resize Increment (MB): 32
Redundancy: mirror
Stripe Columns: 4
Stripe Width (KB): 128
Disk Group: DATA
Volume: VOLUME1
Compatible.advm : 11.2.0.0.0

$ /sbin/advmutil volinfo -l /dev/asm/volume1-228
Device : /dev/asm/volume1-228 : Interface Version : 1 : Size (MB) : 256 :
Resize Increment (MB) : 32 : Redundancy : mirror : Stripe Columns : 4 :
Stripe Width (KB) : 128 : Disk Group : DATA : Volume : VOLUME1 :
Compatible.advm : 11.2.0.0.0

$ /sbin/advmutil volinfo -L /dev/asm/volume1-228
/dev/asm/volume1-228 1 256 32 mirror 4 128 DATA VOLUME1 11.2.0.0.0
```


Part IV

Advanced Topics

Part IV contains the following chapters:

- [Understanding Oracle ACFS Advanced Topics](#) (page 17-1)
- [Managing Oracle Flex ASM](#) (page 19-1)
- [Administering Oracle ASM Filter Driver](#) (page 20-1)
- [Creating Database Files in Oracle ASM Using SQL*Plus](#) (page 21-1)

Understanding Oracle ACFS Advanced Topics

Oracle ACFS advanced topics include discussions about more complex administrative issues.

This appendix discusses Oracle Automatic Storage Management Cluster File System (Oracle ACFS) advanced topics, including limits, advanced administration, and troubleshooting.

See Also:

Articles available at My Oracle Support (<https://support.oracle.com>) for information about Oracle ACFS and Oracle ADVM.

This appendix contains the following topics:

- [Limits of Oracle ACFS](#) (page 17-1)
- [Limits of Oracle ADVM](#) (page 17-4)
- [Oracle ACFS Loopback Support](#) (page 17-4)
- [Oracle ACFS Drivers Resource Management](#) (page 17-4)
- [Oracle ACFS Registry Resource Management](#) (page 17-5)
- [Oracle ACFS File System Resource Management](#) (page 17-5)
- [Oracle ACFS and Oracle Restart](#) (page 17-6)
- [Oracle ACFS Driver Commands](#) (page 17-6)
- [Oracle ACFS Plug-in Generic Application Programming Interface](#) (page 17-8)
- [Oracle ACFS Tagging Generic Application Programming Interface](#) (page 17-12)
- [Oracle ACFS Diagnostic Commands](#) (page 17-18)
- [Understanding Oracle ACFS I/O Failure Console Messages](#) (page 17-29)

For an overview of Oracle ACFS, see [Introducing Oracle ACFS and Oracle ADVM](#) (page 11-1).

17.1 Limits of Oracle ACFS

The limits of Oracle ACFS are discussed in this section.

The topics contained in this section are:

- [Oracle ACFS Disk Space Usage](#) (page 17-2)
- [Oracle ACFS Error Handling](#) (page 17-3)
- [Oracle ACFS and NFS](#) (page 17-3)

Note:

Oracle ACFS does not support hard links on directories.

17.1.1 Oracle ACFS Disk Space Usage

Oracle ACFS supports 64 mounted file systems on 32-bit systems, and 256 mounts on 64-bit systems. However, more file systems can be mounted if there is adequate memory.

Oracle ACFS supports 2^{40} (1 trillion) files in a file system. More than 4 billion files have been tested. There is no absolute limit to the number of directories in a file system; the limit is based on hardware resources.

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 megabytes 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 Linux `df` command even though some space may not actually be allocated yet. This local storage pool can be as large as 128 megabytes per node and can allow space allocations to succeed, even though commands, such as `df`, report less space available than what is being allocated.

For Oracle ACFS 12.2.0.1, the maximum size that can be allocated to an Oracle ACFS file system is 128 terabytes (TB), as shown in [Table 17-1](#) (page 17-2). The storage limits for Oracle ACFS and Oracle ASM are dependent on disk group compatibility attributes.

Table 17-1 *Maximum file sizes for Oracle ADVM volumes*

Redundancy	Disk Group with COMPATIBLE.ASM < 12.2.0.1	Disk Group with COMPATIBLE.ASM = 12.2.0.1
External	128 TB	128 TB
Normal	64 TB	128 TB
High	42.6 TB	128 TB

See Also:

- [Considerations When Setting Disk Group Compatibility](#) (page 4-60) for information about file size limits and disk group compatibility settings
 - [Oracle ASM Storage Limits](#) (page 4-49) for information about storage limits for Oracle ASM files and disk groups
 - [Table 4-6](#) (page 4-60) for information about Oracle ASM file size limits
-

17.1.2 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 the volume enabled followed by a remount of the file system. See "[Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#) (page 16-23)".

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 run 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 a specific 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 Windows `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.

17.1.3 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. You can use any 32-bit number 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 restarts of the same node or across different nodes in the cluster.

When using High Availability NFS for Grid Home Clusters (HANFS), HANFS automatically handles the situation described in the previous paragraph. For information about HANFS, refer to "[High Availability Network File Storage for Oracle Grid Infrastructure](#) (page 11-44)".

17.2 Limits of Oracle ADVM

The limits of Oracle ADVM and these discussed in this section.

The default configuration for an Oracle ADVM volume is 8 columns and a 1 MB stripe width. The volume extent size is 8 megabytes (MB) if the disk group allocation unit (AU) is less than or equal to 8 MB. If the AU size is greater than 8 MB, then the Oracle ADVM volume extent size is equivalent to the disk group AU size.

Oracle ADVM writes data as 1 megabyte (MB) stripe chunks in round robin fashion to each column and fills a stripe set of eight 8 MB extents with 64 stripe chunks before moving to a second stripe set of eight 8 MB extents for volumes greater than 64 MB. 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.

For information about ASMCMD commands to manage Oracle ADVM volumes, refer to [Managing Oracle ADVM with ASMCMD](#) (page 15-1).

17.3 Oracle ACFS Loopback Support

Oracle ACFS supports loopback functionality on the Linux operating system, enabling Oracle ACFS files to be accessed as devices.

An Oracle ACFS loopback device is an operating system pseudo-device that enables an Oracle ACFS file to be accessed as a block device. This functionality can be used with Oracle Virtual Machines (OVM) in support of OVM images, templates, and virtual disks (vdisks) created in Oracle ACFS file systems and presented through Oracle ACFS loopback devices.

Oracle ACFS loopback functionality provides performance gains over NFS. Files can be sparse or non-sparse.

In addition to general loopback support, Oracle ACFS also provides support for loopback direct I/O (DIO) on sparse images.

17.4 Oracle ACFS Drivers Resource Management

Oracle ACFS, Oracle ADVM, and OKS drivers are loaded during the start of the Oracle Grid Infrastructure stack, except in an Oracle Restart configuration. The drivers

remain loaded until the system is rebooted, at which point, they are loaded again when the Oracle Grid Infrastructure stack restarts.

For information about commands to manage Oracle ACFS, Oracle ADVM, and OKS drivers, refer to "[Oracle ACFS Driver Commands](#) (page 17-6)".

17.5 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](#) (page 17-6)".

With Oracle ASM 12c Release 1 (12.1), the Oracle ACFS registry uses the standard single file system resource available through the SRVCTL file system interface. For more information, refer to "[Oracle ACFS File System Resource Management](#) (page 17-5)". Using SRVCTL enables applications to depend on registered file systems, such as for management of the registered file systems using `srvctl filesystem`. By default, `acfsutil registry` shows only file systems that are set to be always mounted, with the `AUTO_START` attribute set to `always`.

The Oracle ACFS registry requires root privileges to register and delete file systems, however, other users can be entitled to start and stop (mount and unmount) the file systems by use of the `user` option.

17.6 Oracle ACFS File System Resource Management

The Oracle ACFS 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](#) (page 17-6)".

Oracle ASM Configuration Assistant (ASMCA) facilitates the creation of Oracle ACFS file system resources (`ora.diskgroup.volume.acfs`). During database creation with Database Configuration Assistant (DBCA), the Oracle ACFS 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 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 due to the start action of the database application.

The start action for an Oracle ACFS file system resource is to mount the file system. This Oracle ACFS 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 Oracle ACFS 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 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 Oracle ACFS file system resources maintains their correct resource state.

17.7 Oracle ACFS and Oracle Restart

Oracle Restart does not support root-based Oracle ACFS resources for this release. Consequently, the following operations are not automatically performed:

- Loading Oracle ACFS drivers
You can load or unload the drivers manually with the `acfsload` command. For more information, refer to "[acfsload](#) (page 17-6)".
- Mounting Oracle ACFS file systems listed in the Oracle ACFS mount registry
You can mount Oracle ACFS file systems manually with the `mount` command. For information, refer to [Managing Oracle ACFS with Command-Line Tools](#) (page 16-1).
- 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.

17.8 Oracle ACFS Driver Commands

This section describes the Oracle ACFS driver commands that are used during installation to manage Oracle ACFS, Oracle ADVM, and Oracle Kernel Services Driver (OKS) drivers. These commands are located in the `/bin` directory of the Oracle Grid Infrastructure home.

- [acfsload](#) (page 17-6)
- [acfsdriverstate](#) (page 17-7)

17.8.1 acfsload

Purpose

`acfsload` loads or unloads Oracle ACFS, Oracle ADVM, and Oracle Kernel Services Driver (OKS) drivers.

Syntax

```
acfsload { start | stop } [ -s ]
```

[Table 17-2](#) (page 17-7) contains the options available with the `acfsload` command.

Table 17-2 Options for the `acfsload` command

Option	Description
<code>start</code>	Loads the Oracle ACFS, Oracle ADVM, and OKS drivers.
<code>stop</code>	Unloads the Oracle ACFS, Oracle ADVM, and OKS drivers.
<code>-s</code>	Operate in silent mode.

Description

You can use `acfsload` to manually load or unload the Oracle ACFS, Oracle ADVM, and OKS drivers.

Before unloading drivers with the `stop` option, you must dismount Oracle ACFS file systems and shut down Oracle ASM. For information about dismounting Oracle ACFS file systems, refer to "[Deregistering, Dismounting, and Disabling Volumes and Oracle ACFS File Systems](#) (page 16-23)".

root or administrator privilege is required to run `acfsload`.

Examples

The following is an example of the use of `acfsload` to stop (unload) all drivers.

```
# acfsload stop
```

17.8.2 acfsdriverstate**Purpose**

`acfsdriverstate` provides information on the current state of the Oracle ACFS, Oracle ADVM, and Oracle Kernel Services Driver (OKS) drivers.

Syntax

```
acfsdriverstate [-orahome ORACLE_HOME ]
                { installed | loaded | version | supported } [-s]
```

[Table 17-3](#) (page 17-7) contains the options available with the `acfsdriverstate` command.

Table 17-3 Options for the `acfsdriverstate` command

Option	Description
<code>-orahome ORACLE_HOME</code>	Specifies the Oracle Grid Infrastructure home in which the user has permission to execute the <code>acfsdriverstate</code> command.
<code>installed</code>	Determines whether Oracle ACFS is installed on the system.
<code>loaded</code>	Determines whether the Oracle ADVM, Oracle ACFS, and OKS drivers are loaded in memory.
<code>version</code>	Reports the currently installed version of the Oracle ACFS system software.

Table 17-3 (Cont.) Options for the `acfsdriverstate` command

Option	Description
<code>supported</code>	Reports whether the system is a supported kernel for Oracle ACFS.
<code>-s</code>	Specifies silent mode when running the command.

Description

You can use `acfsdriverstate` to display detailed information on the current state of the Oracle ACFS, Oracle ADVN, and OKS drivers.

Examples

The following is an example of the use of `acfsdriverstate`.

```
$ acfsdriverstate version
ACFS-9325:    Driver OS kernel version = 2.6.39-400.3.0.el6uek.x86_64(x86_64).
ACFS-9326:    Driver Oracle version = 150728.
ACFS-9212:    Driver build version = MAIN.
```

17.9 Oracle ACFS Plug-in Generic Application Programming Interface

Oracle ACFS plug-in operations are supported through a common, operating system (OS) independent file plug-in (C library) application programming interface (API).

The topics contained in this section are:

- [Oracle ACFS Pre-defined Metric Types](#) (page 17-8)
- [Oracle ACFS Plug-in APIs](#) (page 17-9)

For more information about Oracle ACFS plug-ins, refer to "[Oracle ACFS Plugins](#) (page 11-43)".

17.9.1 Oracle ACFS Pre-defined Metric Types

Oracle ACFS provides the `ACFSMETRIC1_T` and `ACFSMETRIC2_T` pre-defined metric types.

The `ACFSMETRIC1_T` metric set is defined for the storage virtualization model. The metrics are maintained as a summary record for either a selected set of tagged files or all files in the file system. Oracle ACFS file metrics include: number of reads, number of writes, average read size, average write size, minimum and maximum read size, minimum and maximum write size, and read cache (VM page cache) hits and misses.

Example:

```
typedef struct _ACFS_METRIC1 {
    ub2    acfs_version;
    ub2    acfs_type;
    ub4    acfs_seqno;
    ub8    acfs_nreads;
    ub8    acfs_nwrites;
    ub8    acfs_rcachehits;
    ub4    acfs_avgrsize;
    ub4    acfs_avgwsize;
    ub4    acfs_minrsize;
```

```

    ub4      acfs_maxrsize;
    ub4      acfs_minwsize;
    ub4      acfs_maxwsize;
    ub4      acfs_rbytes_per_sec;
    ub4      acfs_wbytes_per_sec;
    ub8      acfs_timestamp;
    ub8      acfs_elapsed_secs;
} ACFS_METRIC1;

```

The `ACFSMETRIC2_T` is a list of Oracle ACFS write description records containing the `fileID`, starting offset, size, and sequence number of each write. The sequence number preserves the Oracle ACFS write record order as preserved by the plug-in driver. The sequence number provides a way for applications to order multiple message buffers returned from the API. It also provides detection of dropped write records due to the application not draining the message buffers fast enough through the API.

The write records are contained within multiple in-memory arrays. Each array of records may be fetched with the API with a buffer size currently set to 1 M. At the beginning of the fetched `ioctl` buffer is a `struct` which describes the array, including the number of records it contains. The kernel buffers drop the oldest write records if the buffers are filled because the buffers are not being read quickly enough.

Example:

```

typedef struct _ACFS_METRIC2 {
    ub2      acfs_version;
    ub2      acfs_type;
    ub4      acfs_num_recs;
    ub8      acfs_timestamp;
    ACFS_METRIC2_REC acfs_recs[1];
} ACFS_METRIC2;

typedef struct _ACFS_FILE_ID {
    ub8      acfs_fenum;
    ub4      acfs_genum;
    ub4      acfs_reserved1;
}

typedef struct _ACFS_METRIC2_REC {
    ACFS_FILE_ID acfs_file_id;
    ub8      acfs_start_offset;
    ub8      acfs_size;
    ub8      acfs_seq_num;
} ACFS_METRIC2_rec;

```

17.9.2 Oracle ACFS Plug-in APIs

Purpose

The Oracle ACFS plug-in application programming interface (API) sends and receives messages to and from the local plug-in enabled Oracle ACFS driver from the application plug-in module.

Syntax

```

sb8 acfsplugin_metrics(ub4 metric_type,
    ub1 *metrics,
    ub4 metric_buf_len,
    oratext *mountp );

```

```
sb8 acfsfileid_lookup(ACFS_FILEID file_id,
    oratext *full_path,
    oratext *mountp );
```

Description

The `acfsplugin_metrics` API is used by an Oracle ACFS application plug-in module to retrieve metrics from the Oracle ACFS driver. The Oracle ACFS driver must first be enabled for plug-in communication using the `acfsutil plugin enable` command. The selected application plug-in metric type model must match the plug-in configuration defined with the Oracle ACFS plug-in enable command. For information about the `acfsutil plugin enable` command, refer to "[acfsutil plugin enable](#) (page 16-156)". The application must provide a buffer large enough to store the metric structures described in "[Oracle ACFS Pre-defined Metric Types](#) (page 17-8)".

If the provided buffer is `NULL` and `metric_buf_len = 0`, the return value is the size required to hold all the currently collected metrics. The application can first query Oracle ACFS to see how big a buffer is required, then allocate a buffer of the necessary size to pass back to Oracle ACFS.

The mount path must be provided to the API to identify the plug-in enabled Oracle ACFS file system that is being referenced.

A nonnegative value is returned for success: 0 for success with no more metrics to collect, 1 to indicate that more metrics are available, or 2 to indicate that no new metrics were collected during the interval. In the case of an error, a negative value is returned and `errno` is set on Linux environments or `SetLastError` is called on Windows.

When using metric type #2, the returned metrics include an `ACFS_FILE_ID`, which contains the fenum and genum pair. In order to translate from the fenum and genum pair to a file path, the application can use `acfsfileid_lookup`. The application must provide a buffer of length `ACFS_FILEID_MAX_PATH_LEN` to hold the path. If there are multiple hard links to a file, the returned path is the first one. This is same behavior when using `acfsutil info id`.

System administrator or Oracle ASM administrator privileges are required to send and receive messages to and from the plug-in enabled Oracle ACFS file system driver.

Writing Applications

To use the plugin API, applications must include the C header file `acfslib.h` which defines the API functions and structures.

```
#include <acfslib.h>
```

When building the application executable, the application must be linked with the `acfs12` library. Check the platform-specific documentation for information about environment variables that must be defined. For example:

```
export LD_LIBRARY_PATH=${ORACLE_HOME}/lib:$
{LD_LIBRARY_PATH}
```

Then when linking, add the `-lacfs12` flag.

Examples

In [Example 17-1](#) (page 17-11), the command enables an Oracle ACFS file system mounted on `/humanresources` for the plug-in service.

Example 17-1 Application Plug-in for Storage Visibility: Poll Model

```
$ /sbin/acfsutil plugin enable -m acfsmetric1 -t HRDATA /humanresources
```

With this command, the application plug-in polls the Oracle ACFS plug-in enabled driver for summary metrics associated with files tagged with HRDATA. The application code includes the following:

```
#include <acfslib.h>
...
/* allocate message buffers */
ACFS_METRIC1 *metrics = malloc (sizeof(ACFS_METRIC1));
/* poll for metric1 data */
while (condition) {
    /* read next summary message from ACFS driver */
    if ((rc = acfsplugin_metrics(ACFS_METRIC_TYPE1, (ub1*)metrics, sizeof(*metrics),
        mountp)) < 0) {
        perror("...Receive failure ... ");
        break;
    }
    /* print message data */
    printf ("reads %8llu ", metrics->acfs_nreads);
    printf("writes %8llu ", metrics->acfs_nwrites);
    printf("avg read size %8u ", metrics->acfs_avgrsize);
    printf("avg write size %8u ", metrics->acfs_avgwsize);
    printf("min read size %8u ", metrics->acfs_minrsize);
    printf("max read size %8u ", metrics->acfs_maxrsize);
    ...
    sleep (timebeforenextpoll);
}
```

In [Example 17-2](#) (page 17-11), the command enables an Oracle ACFS file system mounted on /humanresources for the plug-in service.

Example 17-2 Application Plug-in for File Content: Post Model

```
$ /sbin/acfsutil plugin enable -m acfsmetric1 -t HRDATA -i 5m /humanresources
```

With this command, every 5 minutes the Oracle ACFS plug-in enabled driver posts file content metrics associated with files tagged with HRDATA. In the application code, the call to `acfsplugin_metrics()` is blocked until the metrics are posted. The application code includes the following:

```
#include <acfslib.h>
...
ACFS_METRIC1 *metrics = malloc (sizeof(ACFS_METRIC1));

/* Wait for metric Data */
while (condition) {
    /* Wait for next file content posting from ACFS driver */
    rc = ACFS_PLUGIN_MORE_AVAIL;
    /* A return code of 1 indicates that more metrics are available
    * in the current set of metrics.
    */
    while( rc == ACFS_PLUGIN_MORE_AVAIL) {
        /* This call blocks until metrics are available. */
        rc = acfsplugin_metrics(ACFS_METRIC_TYPE1, (ub1*)metrics, sizeof(*metrics),
            mountp);
        if (rc < 0) {
            perror("...Receive failure ... ");
            break;
        } else if (rc == ACFS_PLUGIN_NO_NEW_METRICS) {
```

```

        printf("No new metrics available.");
        break;
    }
    if (last_seqno != metrics->acfs_seqno-1 ) {
        printf("Warning: Unable to keep up with metrics collection.");
        printf("Missed %d sets of posted metrics.",
            (metrics->acfs_seqno-1)-last_seqno);
    }

    /* print message data */
    printf ("reads %8llu ", metrics->acfs_nreads);
    printf("writes %8llu ", metrics->acfs_nwrites);
    printf("avg read size %8u ", metrics->acfs_avgrsize);
    printf("avg write size %8u ", metrics->acfs_avgwsize);
    printf("min read size %8u ", metrics->acfs_minrsize);
    printf("max read size %8u ", metrics->acfs_maxrsize);
    ...

    last_seqno = metrics->acfs_seqno;
}
}

free(metrics);

```

Example 17-3 Application for Resolving the File Path from a Fenum and Genum Pair

The following example shows how an application resolves a file path from an Oracle ACFS Metric Type 2 record's fenum and genum pair.

```

#include <acfslib.h>
. . .
ACFS_FILE_ID file_id;
char path[ACFS_FILEID_MAX_PATH_LEN];
#ifdef WINDOWS
    file_id.acfs_fenum = atoi(argv[1]);
    file_id.acfs_genum = atoi(argv[2]);
#else
    file_id.acfs_fenum = strtoull(argv[1], 0, 0);
    file_id.acfs_genum = strtoul(argv[2], 0, 0);
#endif
mount_point = argv[3];
sts = acfsfileid_lookup(file_id, path, ACFS_FILEID_MAX_PATH_LEN,
    mount_point, 0);
if (sts < 0)
{
    . . .
}
printf("%s\n", path)

```

17.10 Oracle ACFS Tagging Generic Application Programming Interface

Oracle ACFS tagging operations are supported through a common operating system (OS) independent file tag (C library) application programming interface (API).

An Oracle ACFS tagging API demonstration utility is provided. The demo provides instructions to build the utility with a makefile on each supported platform.

On Solaris, Oracle ACFS tagging APIs can set tag names on symbolic link files, but backup and restore utilities do not save the tag names that are explicitly set on the

symbolic link files. Also, symbolic link files lose explicitly set tag names if they have been moved, copied, tarred, or paxed.

The following files are included:

- `$ORACLE_HOME/usm/public/acfslib.h`
- `$ORACLE_HOME/usm/demo/acfstagsdemo.c`
- `$ORACLE_HOME/usm/demo/Makefile`
Linux, Solaris, or AIX makefile for creating the demo utility.
- `$ORACLE_HOME/usm/demo/MAKEFILE`
Windows demo utility MAKEFILE (for use by the `nmake` utility) to create the demo utility.

The topics contained in this section are:

- [Oracle ACFS Tag Name Specifications](#) (page 17-13)
- [Oracle ACFS Tagging Error Values](#) (page 17-13)
- [acfsgettag](#) (page 17-14)
- [acfslisttags](#) (page 17-15)
- [acfsremovetag](#) (page 17-16)
- [acfssettag](#) (page 17-17)

17.10.1 Oracle ACFS Tag Name Specifications

An Oracle ACFS tag name can be from 1 to 32 characters in length and consist of a combination of the following set of characters only:

- uppercase and lowercase alphabetic characters (A-Z, a-z)
- numbers (0-9)
- hyphen (-)
- underscore (_)
- blank (space)

17.10.2 Oracle ACFS Tagging Error Values

The following are the values for Linux, Solaris, or AIX `errno` in case of failure:

- `EINVAL` – The tag name syntax is invalid or too long.
- `ENODATA` – The tag name does not exist for this file or directory.
- `ERANGE` – The value buffer is too small to hold the returned value.
- `EACCES` – Search permission denied for a directory in the path prefix of path; or the user does not have permission on the file to read tag names.
- `ENAMETOOLONG` – The file name is too long.

- ENOENT – A component of path does not exist.

The following are the values that Windows `GetLastError()` returns in case of failure:

- `ERROR_INVALID_PARAMETER` - The tag name does not exist for this file or directory or the tag name is syntactically incorrect.
- `ERROR_INSUFFICIENT_BUFFER` - The value buffer is too small to hold the returned value.
- `ERROR_ACCESS_DENIED` - Search permission denied for a directory in the path prefix of path; or the user does not have permission on the file to read tag names.
- `ERROR_INVALID_NAME` - The file name or path name is too long or is malformed.
- `ERROR_FILE_NOT_FOUND` - The system cannot find the file specified.

17.10.3 acfsgettag

Purpose

Retrieves the value associated with an Oracle ACFS file tag name.

Syntax

```
sb8 acfsgettag(const oratext *path, const oratext *tagname, oratext *value,
              size_t size, ub4 flags);
```

[Table 17-4](#) (page 17-14) contains the options available with the `acfsgettag` command.

Table 17-4 Options for the `acfsgettag` command

Option	Description
<i>path</i>	Specifies a pointer to a file or directory path name.
<i>tagname</i>	Specifies a pointer to a NULL-terminated Oracle ACFS tag name in the format of a valid tag name for regular files and directories.
<i>value</i>	Specifies the memory buffer to retrieve the Oracle ACFS tag value.
<i>size</i>	Specifies the byte size of the memory buffer that holds the returned Oracle ACFS tag value.
<i>flags</i>	Reserved for future use. Must be set to 0.

Description

The `acfsgettag` library call retrieves the value string of the Oracle ACFS tag name. The return value is the nonzero byte length of the output *value* string on success or `ACFS_TAG_FAIL` on failure. For information about operating system-specific extended error information values that may be obtained when an `ACFS_TAG_FAIL` is returned, refer to "[Oracle ACFS Tagging Error Values](#) (page 17-13)".

Because Oracle ACFS tag names currently use a fixed value string of 0 (the number zero character with a byte length of one) the value is the same for all Oracle ACFS tag name entries. The size of the *value* buffer can be determined by calling `acfsgettag` with a NULL *value* and 0 *size*. The library call returns the byte size necessary to hold the value string of the tag name. `acfsgettag` returns an `ENODATA` error when the tag name is not set on the file.

Examples

[Example 17-4](#) (page 17-15) is an example of the use of the `acfsgettag` function call.

Example 17-4 Retrieving a file tag value

```

sb8 rc;
size_t size;
oratext value[2];
const oratext *path = "/mnt/dir1/dir2/file2";
const oratext *tagname = "patch_set_11_1";
size = 1; (byte)
memset((void *)value, 0, 2*sizeof(oratext));
rc = acfsgettag (path, tagname, value, size, 0);
If (rc == ACFS_TAG_FAIL)
    /* check errno or GetLastError() to process error returns */

```

17.10.4 acfslisttags

Purpose

Lists the tag names assigned to an Oracle ACFS file. For additional information, refer to "[acfsutil tag info](#) (page 16-60)".

Syntax

```
sb8 acfslisttags(const oratext *path, oratext *list, size_t size, ub4 flags);
```

[Table 17-4](#) (page 17-14) contains the options available with the `acfslisttags` command.

Table 17-5 Options for the `acfslisttags` command

Option	Description
<i>path</i>	Specifies a pointer to a file or directory path name.
<i>list</i>	Specifies a pointer to a memory buffer containing the list of Oracle ACFS tag names.
<i>size</i>	Specifies the size (bytes) of the memory buffer that holds the returned Oracle ACFS tag name list.
<i>flags</i>	Reserved for future use. Must be set to 0.

Description

The `acfslisttags` library call retrieves all the tag names assigned to an Oracle ACFS file. `acfslisttags` returns a list of tag names into the *list* memory buffer. Each tag name in the list is terminated with a NULL. If a file has no tag names then the list is empty. The memory buffer must be large enough to hold all of the tag names assigned to an Oracle ACFS file.

An application must allocate a buffer and specify a list size large enough to hold all of the tag names assigned to an Oracle ACFS file. An application can optionally obtain the list buffer size needed by first calling `acfslisttags` with a zero value buffer size and NULL list buffer. The application then checks for nonzero, positive list size return values to allocate a list buffer and call `acfslisttags` to retrieve the actual tag name list.

On success, the return value is a positive byte size of the tag name list or 0 when the file has no tag names. On failure, the return value is `ACFS_TAG_FAIL`. For information about operating system-specific extended error information values that may be obtained when an `ACFS_TAG_FAIL` is returned, refer to "[Oracle ACFS Tagging Error Values](#) (page 17-13)".

Examples

[Example 17-5](#) (page 17-16) is an example of the use of the `acfslisttags` function call.

Example 17-5 Listing file tags

```
sb8 listsize;
sb8 listsize2;
const oratext *path = "/mnt/dir1/dir2/file2";
oratext *list;
/* Determine size of buffer to store list */
listsize = acfslisttags (path, NULL, 0, 0);
if (listsize == ACFS_TAG_FAIL)
/* retrieve the error code and return */

if (listsize)
{
    list = malloc(listsize)
    /* Retrieve list of tag names */
    listsize2 = acfslisttags (path, list, listsize, 0);
    if (listsize2 == ACFS_TAG_FAIL)
        /* check errno or GetLastError() to process error returns */
    if (listsize2 > 0)
        /* file has a list of tag names to process */
    else
        /* file has no tag names. */
}
else
/* file has no tag names. */
```

17.10.5 acfsremovetag

Purpose

Removes the tag name on an Oracle ACFS file.

Syntax

```
sb8 acfsremovetag(const oratext *path, const oratext *tagname, ub4 flags);
```

[Table 17-6](#) (page 17-17) contains the options available with the `acfsremovetag` command.

Table 17-6 Options for the `acfsremovetag` command

Option	Description
<i>path</i>	Specifies a pointer to a file or directory path name.
<i>tagname</i>	Specifies a pointer to a NULL-terminated Oracle ACFS tag name in the format of a valid tag name for regular files and directories.
<i>flags</i>	Reserved for future use. Must be set to 0.

Description

The `acfsremovetag` library call removes a tag name on an Oracle ACFS file. The return value is `ACFS_TAG_SUCCESS` or `ACFS_TAG_FAIL`. For information about operating system-specific extended error information values that may be obtained when an `ACFS_TAG_FAIL` is returned, refer to "[Oracle ACFS Tagging Error Values](#) (page 17-13)".

Examples

[Example 17-6](#) (page 17-17) is an example of the use of the `acfsremovetag` function call.

Example 17-6 Removing file tags

```

sb8 rc;
const oratext *path= "/mnt/dir1/dir2/file2";
const oratext *tagname = "patch_set_11_1";
rc = acfsremovetag (path, tagname, 0);
If (rc == ACFS_TAG_FAIL)
    /* check errno or GetLastError() to process error returns */

```

17.10.6 `acfssettag`

Purpose

Sets the tag name on an Oracle ACFS file. For additional information, refer to "[acfsutil tag set](#) (page 16-61)".

Syntax

```

sb8 acfssettag(const oratext *path, const oratext *tagname, oratext *value,
              size_t size, ub4 flags);

```

[Table 17-7](#) (page 17-17) contains the options available with the `acfssettag` command.

Table 17-7 Options for the `acfssettag` command

Option	Description
<i>path</i>	Specifies a pointer to a file or directory path name.
<i>tagname</i>	Specifies a pointer to a NULL-terminated Oracle ACFS tag name in the format of a valid tag name for regular files and directories.

Table 17-7 (Cont.) Options for the `acfssettag` command

Option	Description
<i>value</i>	Specifies the memory buffer to set the Oracle ACFS tag value.
<i>size</i>	Specifies the byte size of the Oracle ACFS tag value.
<i>flags</i>	Reserved for future use. Must be set to 0.

Description

The `acfssettag` library call sets a tag name on an Oracle ACFS file. The return value is `ACFS_TAG_SUCCESS` or `ACFS_TAG_FAIL`. For information about operating system-specific extended error information values that may be obtained when an `ACFS_TAG_FAIL` is returned, refer to ["Oracle ACFS Tagging Error Values \(page 17-13\)"](#).

Because Oracle ACFS tag names currently use a fixed value string of 0 (the number zero character with a byte length of one) the *value* is the same for all Oracle ACFS tag name entries.

Examples

[Example 17-7](#) (page 17-18) is an example of the use of the `acfssettag` function call.

Example 17-7 Setting file tags

```

sb8 rc;
size_t size;
const oratext *value ;
const oratext *path= "/mnt/dir1/dir2/file2";
const oratext *tagname = "patch_set_11_1";
value = "0"; /* zero */
size = 1; (byte)
rc = acfssettag (path, tagname, (oratext *)value, size, 0);
If (rc == ACFS_TAG_FAIL)
    /* check errno and GetLastError() to process error returns */

```

17.11 Oracle ACFS Diagnostic Commands

This topic provides a summary of the Oracle ACFS command-line utilities for diagnostic purposes.

Oracle ACFS provides various `acfsutil` command-line utilities for diagnostic purposes.

Note:

Run the diagnostic commands only when Oracle Support requests diagnostic data for analysis.

The following table lists the Oracle ACFS utilities with brief descriptions.

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Table 17-8 Summary of Oracle ACFS diagnostic commands

Command	Description
acfsdbg (page 17-19)	Debugs an Oracle ACFS file system.
acfsutil blog (page 17-22)	Writes text to a blog file.
acfsutil dumpstate (page 17-23)	Collects internal Oracle ACFS state information.
acfsutil log (page 17-24)	Retrieves memory diagnostic log files and manages debug settings.
acfsutil meta (page 17-26)	Copies metadata from an Oracle ACFS file system into a separate output file.
acfsutil plogconfig (page 17-28)	Manages Oracle ACFS persistent logging configuration settings.

17.11.1 acfsdbg

Purpose

Debugs an Oracle ACFS file system.

Syntax and Description

```
acfsdbg [-r] [-l] [-x] volume_device
acfsdbg -h
```

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

[Table 17-9](#) (page 17-19) contains the options available with the `acfsdbg` command.

Table 17-9 Options for the acfsdbg command

Option	Description
-h	Prints out the usage message which displays the various options that are available when invoking the <code>acfsdbg</code> command, then exits.
-r	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 because some data is cached by the file system mounts.
-l	Processes kernel log files. The default is to not process the log files.
-x <i>file_name</i>	Specified for accelerator data collected by <code>acfsutil meta</code> . Only used for this type of data.
<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 views and modifies on disk structures of the file system.

Note:

The `acfsdbg` command is intended for experienced development and support engineers to examine on disk structures to diagnose problems. Use with caution.

When `acfsdbg` is started, it displays a command prompt. At the command prompt, you can enter the subcommands listed in [Table 17-10](#) (page 17-20).

You can also use `acfsdbg` for 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 17-10](#) (page 17-20) lists the subcommands of `acfsdbg`.

Table 17-10 Subcommands for `acfsdbg`

Option	Description	Syntax
<code>calculate</code>	Calculates simple 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 the 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>

Table 17-10 (Cont.) Subcommands for acfsdbg

Option	Description	Syntax
fenum	Displays the specified File Entry Table (FETA) entry	<p>fenum [-f -e -d] <i>FETA_entry_number</i></p> <p>-f Displays all on disk structures related to this structure</p> <p>-e Displays all on disk extent information related to this structure</p> <p>-d Casts the structure as a directory and displays its contents</p> <p><i>FETA_entry_number</i> The File Entry Table number used to identify a file on the file system</p>
help	Displays help message	help
offset	Displays structure at disk offset	<p>offset [-c <i>cast</i>] [-f -d] <i>disk_offset</i></p> <p>-f Displays all on disk structures related to this structure</p> <p>-d Casts the structure as a directory and displays its contents</p> <p><i>disk_offset</i> Disk offset to display. The value can be an expression as used by the calculate subcommand</p>
open	Opens a handle to a device. The default is the volume device name entered on the command line	open [<i>volume_device</i>]
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	<p>Reads value from offset</p> <p>The default size to read in is 8 bytes</p> <p>The default count to read is 1</p>	<p>read [-1 -2 -4 -8 -s] [<i>count</i>] <i>offset</i></p> <p>-1 Read byte value</p> <p>-2 Read 2 byte (short) value</p> <p>-4 Read 4 byte (int) value</p> <p>-8 Read 8 byte (long) value</p> <p>-s Read null-terminated string</p> <p><i>count</i> Number of values to read. If not specified, the default is 1</p> <p><i>offset</i> Disk offset to read. The value can be an expression as used by the calculate subcommand</p>
snapshot	Sets the context of commands to the specified snapshot	snapshot <i>snapshot_name</i>

Table 17-10 (Cont.) Subcommands for `acfsdbg`

Option	Description	Syntax
<code>write</code>	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 <code>calculate</code> subcommand Numeric values can also be an expression as used by the <code>calculate</code> subcommand This command is disabled in read-only mode	<code>write [-1 -2 -4 -8 -c -s] [-C -CE] <i>offset value</i></code> -1 Write byte value -2 Write 2 byte (short) value -4 Write 4 byte (int) value -8 Write 8 byte (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 <code>calculate</code> subcommand <i>value</i> The value to write. If numeric, the value can be an expression as used by the <code>calculate</code> subcommand

Examples

[Example 17-8](#) (page 17-22) shows the use of the `acfsdbg` subcommand.

Example 17-8 Using the `acfsdbg` command

```
$ /sbin/acfsdbg /dev/asm/voume1-123
acfsdbg: version = 11.2.0.3.0
Oracle ASM Cluster File System (ACFS) On-Disk Structure Version: 39.0
The ACFS volume was created at Mon Mar 2 14:57:45 2011
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
```

17.11.2 `acfsutil` blog

Purpose

Writes text to the blog file.

Syntax and Description

```
acfsutil [-h] blog
```

```
acfsutil blog {-t text | -u} mount_point
```

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

The following table contains the options available with the `acfsutil blog` command.

Table 17-11 Options for the `acfsutil blog` command

Option	Description
<code>-t text</code>	Writes text to the blog file at the specified mount point.
<code>-u</code>	Updates blog debug levels from <code>dbg</code> file.
<code>mount_point</code>	Specifies the mount point.

The `acfsutil blog` command enables you to write text to a blog file.

.

Examples

The following example illustrates how to run the `acfsutil blog` command. Running `acfsutil blog` with the `-h` option displays help.

Example 17-9 Using `acfsutil blog`

```
$ /sbin/acfsutil -h
```

```
$ /sbin/acfsutil -t "this is a blog test" blog my_mount_point
```

```
$ /sbin/acfsutil -u blog my_mount_point
```

17.11.3 acfsutil dumpstate**Purpose**

Collects internal Oracle ACFS state information for diagnosis by Oracle support.

Syntax and Description

```
acfsutil [-h] dumpstate
```

```
acfsutil dumpstate acfs_path
```

`acfsutil -h dumpstate` displays help text and exits.

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

The following table contains the options available with the `acfsutil dumpstate` command.

Table 17-12 Options for the `acfsutil dumpstate` command

Option	Description
<code>acfs_path</code>	Specifies the directory to a file in an Oracle ACFS file system.

The `acfsutil dumpstate` command collects internal Oracle ACFS state information for a specified file and its corresponding file system. The state information is written to a binary incident file in a logging directory.

Note:

Run the `acfsutil dumpstate` command only when Oracle Support requests diagnostic and debugging data for analysis.

Examples

The following example shows the use of the `acfsutil dumpstate` command.

Example 17-10 Using the `acfsutil dumpstate` command

```
$ /sbin/acfsutil dumpstate /acfsmounts/acfs1/myfile
```

17.11.4 `acfsutil log`

Purpose

Retrieves memory diagnostic log files and manages debug settings.

Syntax and Description

```
acfsutil [-h] log
```

```
acfsutil log [-f filename] [-s] [-r n{K|M|G|T|P}] [-p {avd|ofs|oks}] [-l debuglevel]
            [-n consolelevel] [-o wait_time] [-q] [-c debugcontext] [-T file_type]
            [-m mount_point] [-a] [-C] [-t]
```

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

The following table contains the options available with the `acfsutil log` command.

Table 17-13 Options for the `acfsutil log` command

Option	Description
<code>-f filename</code>	Write the in-memory log to the specified file. The default file is <code>oks.log</code> in the current directory.
<code>-s</code>	Shows the size of the in-memory log file.
<code>-r n{K M G T P}</code>	Sets the size of the in-memory log file.

Table 17-13 (Cont.) Options for the acfsutil log command

Option	Description
-p {avd ofs oks}	Specifies the product for setting the level or querying settings. The is default all products: Oracle ADVM (avd), Oracle ACFS (ofs), and Oracle Kernel Services (oks)
-l <i>debuglevel</i>	Sets the in-memory debug level. The default debug level is 2. Valid values are 0-6.
-n <i>consolelevel</i>	Sets the debug level for persistent logging . Other persistent log configuration settings are managed by the <code>acfsutil plogconfig</code> command.
-o <i>wait_time</i>	Sets the log size, the debug level, and the product values on all nodes; waits for the number of seconds specified by <i>wait_time</i> ; dumps in the memory log on all nodes; and then resets the debug level and the log size.
-q	Queries the debug settings for a specified product. For example: <code>acfsutil log -p avd -q</code>
-c <i>debugcontext</i>	Sets the debug context, internal only.
-T <i>file_type</i>	Sets the debug file type, internal only.
-m <i>mount_point</i>	Specifies to debug only the file system at the specified mount point.
-a	Resets the debug logging to log for all file systems.
-C	Dumps a memory log on all cluster nodes, and also can be added to -t option.
-t	Dumps all Hang Manager thread information to in-memory and persistent logs.

The `acfsutil log` command enables you to manage memory diagnostic log files. With none of the options specified, the `acfsutil log` command retrieves and writes the `./oks.log` memory log by default.

The `-o` option performs the following:

1. Sets the log size to 500M , the log level to 5, and the product to `ofs` (acfs) for the in-memory log on all nodes
2. Displays an informational message, such as *Blocking for 180 seconds, reproduce problem now*
3. After waiting for the specified number of seconds, then displays *Dumping log on all nodes*
4. Initiates a clusterwide dump of logs
5. Resets the log level to 2 and resets the log size to the default

The `-o` option can be combined with the `-p`, `-l`, and `-r` options if the default product, debug level, or log size settings should be changed.

You must be the `root` user or an Oracle ASM administrator user to run this command.

Examples

The following example illustrates how to run the `acfsutil log` command to display help.

Example 17-11 Using acfsutil log

```
$ /sbin/acfsutil -h log
```

17.11.5 acfsutil meta

Purpose

Copies metadata from an Oracle ACFS file system into a separate output file.

Syntax and Description

```
acfsutil meta -h
acfsutil meta [-v] [-q] [-l log_file_path] [-o acfs_extent_offsets] {-f output_file}
[-a <device>] volume_device
```

`acfsutil meta -h` displays help text and exits.

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

[Table 17-14](#) (page 17-26) contains the options available with the `acfsutil meta` command.

Table 17-14 Options for the acfsutil meta command

Option	Description
<code>-v</code>	Specifies verbose mode to generate additional diagnostic messages.
<code>-q</code>	Invokes the metadata collector in <i>quick</i> mode.
<code>-l log_file_path</code>	Specifies the path to the log file. If not specified, the log file is generated in the current directory with a default name of <code>acfs.meta.log</code> .
<code>-o acfs_extent_offsets</code>	Specifies a list of comma separated file offsets from which the meta collector additionally copies data.
<code>-f output_file</code>	Specifies the path name of the output file into which the metadata is copied.
<code>-a device</code>	Specifies the location of any associated accelerator device, to be used if the file system is unmountable.
<code>volume_device</code>	Specifies a volume device name of the file system which is to be copied.

The `acfsutil meta` command operates as a metadata collector to partially copy an Oracle ACFS file system into a separate specified output file. The metadata collector reads the contents of the file system specified by the volume device name of an Oracle ACFS file system. This input file system is searched for Oracle ACFS metadata and then all metadata found is written into the specified output file. The generated output file can be used for further diagnostics and analysis, without impact to the original file system at the customer site.

To obtain the best copy of the file system with `acfsutil meta`, unmount the file system before running `acfsutil meta`. If it is not possible to unmount the file system, avoid modifying the contents or performing a volume resizing operation while `acfsutil meta` is running.

If the original file system is very large, then the output file can also be very large. The output file should be placed on a file system that supports sparse files because this placement can reduce the size of the file. When copying the output file, use a utility that supports sparse files. Compress the output file when possible to reduce storage space and transmission time.

If the file system has an accelerator device associated with it, `acfsutil meta` also copies the accelerator device data to a second output file. The second file uses the file name from the `-f` option with an appended `.acc` extension. For example, if you specify `acfsutil meta -f /tmp/mymetafile volume1-123`, the meta collector places a copy of `volume1-123` in `/tmp/mymetafile`, and the copy of its accelerator device in the `/tmp/mymetafile.acc` file. This operation occurs automatically.

In most circumstances, `acfsutil meta` automatically copies the accelerator device to the second file. However, if you have think that the meta collector is not able to find the accelerator device on its own, you can specify the name on the command line with the `-a` option. For example, this situation could occur if the file system is corrupt. Note that using the `-a` option overrides how the meta collector operates automatically, so `-a` should be used carefully.

The output file should not be placed on the Oracle ACFS device that is specified as the input device because the metadata command might process the output file also. The output file should be placed on a file system that can support an output file which is the size of the Oracle ACFS input volume device.

The `-q` flag should be used with caution. When `-q` is specified, the meta collector does not scan and copy the entire input file system. Instead, it only scans and copies a predetermined number of bytes and certain data structures which are considered important. The primary use for the `-q` flag is for situations where there is not sufficient time to run the full version of the metadata collector. The `-q` flag should not be used unless it is recommended by the support personnel investigating the problem.

The `acfsutil meta` output file can be read by the `acfsutil fsck` command in most cases. However, the Oracle ACFS specific `acfsutil fsck` command on some OS platforms might not access the output file correctly or might not work with a specified flag. You can use a slightly modified `fsck` command form in these cases. For example:

- On Linux, run the command in this format if you are using the `-x` flag:

```
/sbin/fsck.acfs -x accelerator_meta_file meta_output_file
```

- On AIX, run the command in this format:

```
/sbin/helpers/acfs/fsck meta_output_file
```

Examples

[Example 17-12](#) (page 17-28) shows the use of the `acfsutil meta` command.

Example 17-12 Using the `acfsutil meta` command

```
$ /sbin/acfsutil meta -f /acfsmounts/critical_apps/meta_output_file /dev/asm/
volume1-123
```

17.11.6 acfsutil plogconfig

Purpose

Manages Oracle ACFS persistent logging configuration settings.

Syntax and Description

```
acfsutil plogconfig [-h] [-d persistent_log_directory] [-t] [-q ] [-i seconds]
                    [-s buffer_size] [-l low_water_percent] [-u high_water_percent]
                    [-m max_logfile_size] [-n max_logfile_number]
```

`acfsutil -h plogconfig` displays help and exits.

For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

The following table contains the options available with the `acfsutil plogconfig` command.

Table 17-15 Options for the `acfsutil plogconfig` command

Option	Description
<code>-d</code> <i>persistent_log_directory</i>	Specifies an alternative logging directory. If not specified, the default directory is <code>\$ORACLE_BASE/crsdata/hostname/acfs</code> .
<code>-t</code>	Terminates logging.
<code>-q</code>	Queries for and then displays the persistent logging configuration settings.
<code>-i</code> <i>seconds</i>	Specifies the number of seconds for the interval timer.
<code>-s</code> <i>buffer_size</i>	Sets the log buffer size in kilobytes.
<code>-l</code> <i>low_water_percent</i>	Sets the file write trigger as a percentage.
<code>-u</code> <i>high_water_percent</i>	Sets the file write throttle as a percentage.
<code>-m</code> <i>max_logfile_size</i>	Sets the maximum log file size in megabytes.
<code>-n</code> <i>max_logfile_number</i>	Sets the maximum number of log files.

The `acfsutil plogconfig` command provides a diagnostic tool to manage configuration settings for persistent logging.

All command arguments are optional, but at least one argument must be specified.

Note:

Run the `acfsutil plogconfig` command only when Oracle Support requests configuration of persistent logging settings.

You must be the `root` user or an Oracle ASM administrator user to run this command.

Examples

The following example illustrates the use of the `acfsutil plogconfig` command to display the current configuration settings.

Example 17-13 Using the Oracle ACFS `acfsutil plogconfig` command

```
# /sbin/acfsutil plogconfig -q

Log Directory Name : /oracle/crsdata/my_host/acfs
Buffer Size (KB) : 64
Low Water Level (percent) : 50
High Water Level (percent) : 75
Timer Interval (Seconds) : 5
Maximum Number of Log Files : 10
Maximum Log File Size (MB) : 100
```

17.12 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 ADVN 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 finds 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/log/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 .
```

Configuring Oracle ACFS Snapshot-Based Replication

The requirements for Oracle ACFS snapshot-based replication are discussed in this section.

This chapter describes how to configure Oracle ACFS snapshot-based replication available with release 12.2. As with Oracle ACFS replication installations before release 12.2, the overall functional goal of snapshot-based replication is to ensure that updates from a primary cluster are replicated to a standby cluster. However, the snapshot based replication technology uses snapshots of the primary file system and transfers the differences between successive snapshots to the standby file system using the standard `ssh` command. Oracle ACFS replication functionality before release 12.2 replicated changes continuously, building on Oracle networking technologies, notably Network Foundation Technologies (NFT), to ensure connectivity between the primary and standby clusters.

This change in the design and implementation of Oracle ACFS replication introduces some differences in how replication is configured and used. For example, the use of `ssh` requires setting up host and user keys appropriately on the primary and standby nodes where replication is performed.

This chapter contains the following topics:

- [Configuring `ssh` for Use With Oracle ACFS Replication](#) (page 18-1)
- [Installing `ssh` and `Cygwin` on Windows](#) (page 18-5)
- [Upgrading to Oracle ACFS Snapshot-Based Replication](#) (page 18-12)

For an overview of Oracle ACFS replication, refer to [Oracle ACFS Replication](#) (page 11-36). For information about Oracle ACFS replication commands, refer to [Oracle ACFS Command-Line Tools for Replication](#) (page 16-63).

18.1 Configuring `ssh` for Use With Oracle ACFS Replication

This section describes how to configure `ssh` for use Oracle ACFS snapshot-based replication available with release 12.2.

Choosing an Oracle ACFS Replication User

Oracle ACFS snapshot-based replication uses `ssh` as the transport between the primary and standby clusters, so the user identity under which replication is performed on the standby must be carefully managed. In the replication process, the `root` user (or local `SYSTEM` on Windows) on the primary node where replication is running uses `ssh` to log in to the standby node involved in replication.

Because it is not advisable for `ssh` to log in as `root` on the standby node, a minimally-privileged user identity should be used. The user chosen should have Oracle ASM administration privileges. Usually, the user specified to the Oracle installer when the

Oracle software was first installed belongs to the needed groups, so can be convenient to choose as the replication user. In this discussion, the replication user is identified as *repluser*; however, you would replace *repluser* with the actual ASM user name that you have selected. For information about user privileges for Oracle ASM, refer to [About Privileges for Oracle ASM](#) (page 3-28). For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Distributing Keys for Oracle ACFS Replication

Getting a public key for root from the primary cluster

A public key for `root` defined on each node of your primary cluster must be known to *repluser* on each node of your standby cluster.

To make the key known, the directory `~repluser/.ssh` must exist on each standby node. If this directory does not exist, then create it with access only for *repluser*. Ensure that an `ls` command for the `.ssh` directory displays output similar to:

```
repluser@standby $ ls -ld ~/.ssh
drwx----- 2 repluser dba 4096 Jan 27 17:01 .ssh
```

If a public key for `root` is defined on a given primary node, then it resides in a `.pub` file, such as `/root/.ssh/id_rsa.pub`. If a public key file exists, then add its contents to the set of keys authorized to log in as *repluser* on each node of the standby where replication is run. Append the key to the file `~repluser/.ssh/authorized_keys2` on each standby node, creating this file if necessary.

If a public key file does not exist, generate a public and private key pair on the primary by running the following command as `root`.

```
# ssh-keygen -t rsa
```

You can press the enter key in response to each prompt issued by the command. Copy the resulting `.pub` file to each standby node.

You have the option to share the same public/private key pair for `root` across all of the nodes in your primary cluster, or to establish a different key pair for each primary node. If the same public key is valid for `root` across all nodes in your primary cluster, then only that key must be added to the file `~repluser/.ssh/authorized_keys2` on each node of your standby cluster. If each primary node has its own public key for `root`, then all the public keys must be added to the file. In either case, you can minimize work by copying the updated `authorized_keys2` file on a given node of the standby to the other nodes of the cluster.

Getting host keys for the standby cluster

A host key for each standby node where replication may run must be known on each primary node where replication may run. One way to generate the correct key is to run `ssh` manually as `root` from each primary node to each standby node. If the correct host key is not known already, then a warning displays and you can enable `ssh` to add the key.

Note that there are two users involved in the `ssh` connection. While `ssh` on the primary node connects to the standby node as `root`, `ssh` logs in on the standby node as *repluser*. Any command run by `ssh` on the standby runs with the privileges of *repluser*, not with `root` privileges.

Because the primary node connects to the standby node as user `root`, the host key for the standby node must be added to the `known_hosts` file of the `root` user, not the file for *repluser*. The following is an example of obtaining a host key:


```
[root@primary usm]# ssh repluser@standby date
The authenticity of host 'standby (10.137.13.85)' can't be established.
RSA key fingerprint is 1b:a9:c6:68:47:b4:ec:7c:df:3a:f0:2a:6f:cf:a7:0a.
Are you sure you want to continue connecting (yes/no)?
```

If you respond with *yes*, then the ssh setup is complete. A host key for host *standby* is stored in the *known_hosts* file (*~root/.ssh/known_hosts*) on the host *primary* for the user *root*.

After the host key setup for standby nodes is complete on a given primary node, you need to perform an additional step if you use a Virtual IP address (VIP) to communicate with your standby cluster. You must add the VIP name or address at the start of each line of the *known_hosts* file that refers to a host in the standby cluster. For example, if you use a VIP with the name *standby12_vip*, and your *known_hosts* file contains the following two lines that refer to your standby:

```
standby1,10.242.20.22 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQ3pM2YTd4UUUEWEoCKDGgaTgsmPkQToDrdtU+JtVIq/96muvU
BaJUK83aqzeNIQkh+hUULsUdgKoKT5bXrWYqhY6AlTEqNgBHjBrJt9C73BbQd9y48jSc2G+WQWyuI/
+s1Q+hIJdBNMxvMBQafisPWWUcaIx9Y/Jz1PgF61RP2cbfqAzixDot9fqRrAKL3G6A75A/6TbwmEW07d1zqOv
17ZGyeDYf5zQ72F/VOP9UGMeT/5DmcYTn3kTVGjOTbnRBe4A41Y4rVw5c+nZBDFre66XtORfQgwQB5ztW/Pi
08GYbcIszKoZx2HST9AZxYIAgcrnNYG2Ae0K6QLxxxScP
standby2,10.242.20.23 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDIsczjzNtKN03SY8K1846skFTVP1HF/ykswbmktEjL6KTWTW+NR
U4MGbvkBqqdXxuPCR7aogO2U3PEOglUVf3DWUoux8IRvqKU+dJcdTibMFkDAIhTnzb14gZ/LRTjn+GYsuP5
Qz2vgL/U0ki887mZCRjwVl1b5FNH8sXBUV2QcD7bjf98VXF6n4gd5UiIC3jv612nVTKDwtNHpUTS1dQAI+1D
tr0AieZTsuxXMaDdUZHGKDotjciMB3mCkKm/u3IFoioDqdZE4+vITX9G7DBN4CVPXawp+b5Kg8X9P+08Eehu
tM1BJ5lafy1bxoVlXUDLVIIFBJNKrsqBvxxxpS7
```

To enable the use of the VIP, you would modify these two lines to read as follows:

```
standby12_vip,standby1,10.242.20.22 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQ3pM2YTd4UUUEWEoCKDGgaTgsmPkQToDrdtU+JtVIq/96muvU
BaJUK83aqzeNIQkh+hUULsUdgKoKT5bXrWYqhY6AlTEqNgBHjBrJt9C73BbQd9y48jSc2G+WQWyuI/
+s1Q+hIJdBNMxvMBQafisPWWUcaIx9Y/Jz1PgF61RP2cbfqAzixDot9fqRrAKL3G6A75A/6TbwmEW07d1zqOv
17ZGyeDYf5zQ72F/VOP9UGMeT/5DmcYTn3kTVGjOTbnRBe4A41Y4rVw5c+nZBDFre66XtORfQgwQB5ztW/Pi
08GYbcIszKoZx2HST9AZxYIAgcrnNYG2Ae0K6QLxxxScP
standby12_vip,standby2,10.242.20.23 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDIsczjzNtKN03SY8K1846skFTVP1HF/ykswbmktEjL6KTWTW+NR
U4MGbvkBqqdXxuPCR7aogO2U3PEOglUVf3DWUoux8IRvqKU+dJcdTibMFkDAIhTnzb14gZ/LRTjn+GYsuP5
Qz2vgL/U0ki887mZCRjwVl1b5FNH8sXBUV2QcD7bjf98VXF6n4gd5UiIC3jv612nVTKDwtNHpUTS1dQAI+1D
tr0AieZTsuxXMaDdUZHGKDotjciMB3mCkKm/u3IFoioDqdZE4+vITX9G7DBN4CVPXawp+b5Kg8X9P+08Eehu
tM1BJ5lafy1bxoVlXUDLVIIFBJNKrsqBvxxxpS7
```

Ultimately, the host key configuration performed on this first node of your primary cluster must be performed on every node in your primary cluster; the result of the above sequence, or an equivalent, must exist on each primary node. One way to minimize the manual effort required to achieve this configuration is to update the *known_hosts* file on one node of the primary cluster, then copy the updated file to the other nodes of the cluster.

Note:

By default, replication enables strict host key checking by `ssh`, to ensure that the primary node connects to the intended standby node or cluster when it runs `ssh`. However, if you are certain that this checking is unneeded, such as the case when the primary and standby clusters communicate over a private network, the use of strict host key checking by `ssh` can be disabled. For information about disabling strict host key checking, refer to the `-o sshStrictKey=no` option of the `acfsutil repl init primary` command. If strict host key checking is disabled, then no host key setup is required. For information about the `acfsutil repl init` command, refer to [acfsutil repl init](#) (page 16-69).

Notes on permissions for ssh-related files

For `ssh` to work with the keys you have established, you must ensure that permissions are set properly on each node for the relevant `.ssh` directory and some of the files the directory contains. On each primary node, this refers to the `.ssh` directory for `root`. On each standby node, the `.ssh` directory to check is the one for `repluser`.

For details on the permissions that should be given to each `.ssh` directory and key files within the directory, refer to the documentation for your `ssh` implementation, such as the `FILES` section of the `ssh(1)` manual page.

Notes on sshd configuration

After you begin using replication, `ssh` is started frequently to perform replication operations. On some platforms, the `ssh` daemon `sshd` may be configured to log a message through `syslog` or a similar facility each time an `ssh` connection is established. To avoid this, the server configuration file `/etc/ssh/sshd_config` can be modified to specify a lower frequency of logging. The parameter that controls logging is called `LogLevel`. Connection messages are issued at level `INFO`. Any lower `LogLevel` setting, such as `ERROR`, suppresses those messages. For example, you can suppress log messages by adding the following line to the file:

```
LogLevel ERROR
```

Validating your ssh-related key configuration

After you have established the host and user keys for `ssh`, you can use the command `acfsutil repl info -c -u` to validate the keys. You run this command as `root` on each node of the primary cluster. It takes as arguments all the hostnames or addresses on the standby cluster that the primary may use in the future to perform replication.

If you are not using a VIP to connect to your standby cluster, then for a given replication relationship, only one standby hostname or address is provided to `acfsutil repl init primary`. However, if future relationships may involve other standby host addresses, specify the complete set of addresses when running the `acfsutil repl info -c -u` command.

If you are using a VIP to connect to your standby cluster, then you should specify the names or host-specific addresses of all standby hosts on which the VIP may be active. Do not specify the VIP name or an address associated with the VIP. When replication uses `ssh` to connect to a VIP, the host key returned is the key associated with the host where the VIP is currently active. Only the hostnames or addresses of individual standby nodes are used by `ssh` in this situation.

The validation command has the following format:

```
# acfsutil repl info -c -u repluser standby-addr1 [standby-addr2 ...] standby-
mountpoint
```

The command confirms that user *repluser* can use *ssh* to connect to each *standby-addr* given in the same way as replication does when initializing.

To validate the key setup for the *standby1* and *standby2* cluster presented in the previous command, you can use the following command:

```
# acfsutil repl info -c -u repluser standby1 standby2 standby-mountpoint
```

The same command would be used if you were planning to use the VIP *standby12_vip* to connect to the cluster.

If you plan to disable strict host key checking, you can skip this checking by adding the `-o sshStrictKey=no` option to the command line.

18.2 Installing ssh and Cygwin on Windows

This section describes how to install Cygwin and start the *ssh* daemon on Microsoft Windows hosts. This information is only applicable when you want to use Oracle ACFS snapshot-based replication on a Windows host.

Cygwin Requirement for Oracle ACFS Replication

Oracle ACFS snapshot-based replication uses *ssh* to transfer replication-related data from the primary cluster to the standby cluster. When you use replication on a host running on Microsoft Windows, you must install Cygwin and start the *ssh* daemon on the host, as described in this section.

Cygwin is essentially a utility that offers a Linux-like environment on a Microsoft Windows host. Technically, it is a DLL (*cygwin1.dll*) that acts as a Linux API layer providing substantial Linux API functionality. After you install Cygwin, you can configure the *ssh* daemon on the host. Oracle ACFS replication is certified and supported with Cygwin 2.0.

The *ssh* daemon enables Oracle ACFS Replication to establish *ssh* connectivity from the primary cluster running replication to the standby cluster running replication.

Note:

These instructions for installing the *ssh* daemon result in a daemon that is usable only to log in and run commands as the specific user *repluser*. The daemon is not intended to work correctly when logging in and running commands as any other user. This restriction results from two factors:

- Users defined for Oracle ACFS must be domain-based users.
 - Enabling the general-purpose use of the *ssh* daemon for a domain-based user involves creating a domain policy that gives that user some very powerful privileges, including *SeTcbPrivilege* which lets the user act as part of the *Trusted Computing Base*. Creating such a policy would not likely be allowed in most corporate environments, so the use of such a policy is not included in these instructions.
-
-

Steps Before Installing Cygwin

Before starting with the `ssh` setup, ensure you are not using `OpenSSH` and `MKSNT` by performing the checks in the following steps.

Note that the navigational steps described in this section may vary for different Microsoft Windows operating systems.

1. Ensure `OpenSSH\bin` and any directory where `mksnt` executable files are stored are not in your `PATH` environment variable. If they are, remove them from the path string by performing the following:
 - a. Right-click on **My Computer** and go to **Properties**.
 - b. In the **System Properties** window, click **Advanced**.
 - c. In this tab, click **Environment Variables**.
 - d. Search for the `PATH` system variable, select it, and if `OpenSSH\bin` or `mksnt`-related values are present in `PATH` string, click **Edit**.
 - e. In the **Edit System Variable** dialog box, delete these values from `PATH` string, then click **OK**.
2. Stop and disable the `ssh` daemon if it is running from `OpenSSH`, `MKS` or any other source. If the `ssh` daemon is running, stop and disable it by doing the following:
 - a. Right-click on **My Computer**, and select **Manage**.
 - b. In the **Computer Management** window, in the left pane, expand **Services and Applications**, and select **Services**.
 - In the right pane, right-click the `ssh` daemon/`MKS Secure Shell` service, then click the **Stop** entry on the menu that appears.
 - In the same pane, right-click the `ssh` daemon/`MKS Secure Shell` service, then click the **Properties** entry on the menu that appears. In the dialog box that opens, select **Disabled** as the **Startup** type, and click **OK**.

Installing Cygwin

To install Cygwin on a Microsoft Windows host, perform the following steps:

1. Access the following URL, then click **Install Cygwin**:
`http://www.cygwin.com/`
2. Depending on the version of Microsoft Windows you are running, download the 32-bit version or the 64-bit version of the Cygwin setup executable.
3. Run the setup executable, then click **Next** to proceed.
4. On the **Choose Installation Type** screen, select **Install from Internet**, then click **Next**.
5. On the **Choose Installation Directory** screen, enter `C:\cygwin` as the **Root Directory**, then click **Next**.
6. On the **Select Local Package Directory** screen, select a directory on your local machine where you want to store the downloaded installation files, then click **Next**.

7. On the **Select Connection Type** screen, select appropriate settings to connect to the internet, then click **Next**.
8. On the **Choose Download Site(s)** screen, select any site from the available list, then click **Next**.
9. On the select packages screen, ensure that you select the following packages, then click **Next**. From the **Archive** category, select **unzip** and **zip**. From the **Net** category, select **openssh** and **openssl**. After selecting the packages and clicking **Next**, the **Resolving Dependencies** screen is displayed. Click **Next** to proceed.
10. On the **Installation Status and Create Icons** screen, do not make any changes. Click **Finish** to complete the installation process.

Configuring ssh

This section describes how to configure `ssh` for Oracle ACFS replication after installing Cygwin on a host.

Choosing Replication Users

Oracle ACFS Replication uses `ssh` as the transport between the primary and standby clusters, so the user identities under which replication is performed need to be explicitly specified. In the replication process, a privileged user on the primary node uses `ssh` to log in to the standby node involved in replication. It is not desirable for `ssh` to log in as a privileged user on the standby node. Instead, a minimally-privileged user identity should be used. The user chosen must be a domain-based user, and must be a member of the `ORA_ASMADMIN` group.

In this discussion, *repladmin* and *repluser* are used to refer to the replication users; however, you would replace *repladmin* and *repladmin* with the actual user names that you have selected. The *repladmin* user refers to the user that runs the `acfsutil repl init` command on both the primary and the standby. This user must be a member of the `Administrators` group. The *repluser* user refers to the user that `ssh` uses to log in on the standby. This user must be a domain-based user, and must be a member of the `ORA_ASMADMIN` group, but not a member of the `Administrators` group. If a suitable user already exists for the use of Oracle RAC, that user should be used as *repluser*.

For information about user privileges for Oracle ASM, refer to [About Privileges for Oracle ASM](#) (page 3-28). For information about running Oracle ACFS `acfsutil` commands, refer to [About Using Oracle ACFS Command-Line Tools](#) (page 16-2).

Configuration Steps

While configuring `ssh`, you may need to run the `cygwin.bat` script. While running `cygwin.bat` on Microsoft Windows Server 2008 and Microsoft Windows Vista, ensure that you invoke the batch file in administrator mode. To do this, right-click the `cygwin.bat` file and select `Run as administrator`.

To configure `ssh` and test your Cygwin setup, follow these steps:

1. After you install Cygwin, navigate to the `C:\cygwin` directory, open the `Cygwin.bat` file in edit mode using any editor, and add the following line before invoking the bash shell.

```
set CYGWIN=binmode ntsec
```

The following lines are the possible contents of the `Cygwin.bat` file after adding the previous line:

```
@echo off
C:
chdir C:\cygwin\bin
set CYGWIN=binmode ntsec bash --login -i
```

2. To verify if Cygwin (`cygrunsrv`) is installed properly, run `C:\cygwin\Cygwin.bat`, and run the following command:

```
cygrunsrv -h
```

If Cygwin is installed properly, then all Cygwin help options are displayed on the screen. If the command returns an error message, then you may have to reinstall Cygwin.

3. Define the *repladmin* and *repluser* identities at the Windows level.

You can also use existing users instead for these roles. If you define new identities for these roles, note the following:

- The *repladmin* user must be a member of the Administrators group. It is recommended that *repladmin* also be a domain-based user. If *repladmin* is a local user, it should be defined with the same group memberships on all primary nodes.
 - The *repluser* user must be a domain-based user, and must be a member of the `ORA_ASMADMIN` group.
4. Configure the `sshd` service. Run `C:\cygwin\Cygwin.bat`, and execute the following command:

```
ssh-host-config
```

After running the command, you are prompted with the following questions. Appropriate responses are shown in bold. Other output may also appear, but is not shown here.

```
*** Query: Should StrictModes be used? (yes/no) yes
*** Query: Should privilege separation be used? <yes/no>: yes
*** Query: New local account 'sshd'? <yes/no>: yes
*** Query: Do you want to install sshd as a service?
*** Query: <Say "no" if it is already installed as a service> <yes/no>: yes
*** Query: Enter the value of CYGWIN for the daemon: [] binmode ntsec
```

Now `ssh-host-config` outputs some notes concerning the user account required to use passwordless logins. This capability is required by Oracle ACFS replication. The *repluser* account should be specified as this account.

```
*** Info: The following privileged accounts were found: 'cyg_server' .
*** Info: This script plans to use 'cyg_server'.
*** Info: 'cyg_server' will only be used by registered services.
*** Query: Do you want to use a different name? (yes/no) yes
```

You should respond **yes** to the prompt, and should specify *repluser* as the name under which you run `sshd`. You are then prompted with the following questions.

```
*** Query: Enter the new user name: repluser
*** Query: Reenter: repluser

***Warning: The specified account 'repluser' does not have the
***Warning: required permissions or group memberships. This may
***Warning: cause problems if not corrected; continuing...
```

```
*** Query: Please enter the password for user 'repluser':
*** Query: Reenter:
```

The above assumes that the user *repluser* already exists. You may ignore the warning output about missing permissions or group memberships. If the configuration is successful, the following message displays.

```
Host configuration finished. Have fun!
```

5. Ensure that the directory `/var/empty` exists and is owned by *repluser*.
6. Ensure that the *repluser* and *repladmin* users are known to Cygwin.

Backup the `c:\cygwin\etc\passwd` file and then open the file in edit mode. Remove any lines from this file that refer to the *repladmin* or *repluser* users. For each user, run the following command. Both users are assumed to be domain-based users.

```
/bin/mkpasswd -d -u repladmin >> /etc/passwd
```

Ensure that the home directory named for each user in `/etc/passwd` exists. If necessary, create it. For example, if `/home/repladmin` is the directory shown for user *repladmin*, perform the following to create the necessary directory.

```
mkdir -p /home/repladmin
chown repladmin /home/repladmin
```

7. Ensure that the ssh daemon starts, and is configured to start automatically:
 - a. Right-click on **My Computer**, and select **Manage**.
 - b. In the **Computer Management** window, in the left pane, expand **Services and Applications**, and select **Services**.
 - c. In the right pane, right-click the **CYGIN sshd** service, then click the **Properties** entry on the menu that appears. In the dialog box that opens, select **Automatic** as the **Startup** type, and click **OK**.

If the ssh daemon does not start, view the `c:\cygwin\var\log\sshd.log` file for information that relates to the failure of the startup process.

Getting the public key for repladmin from the primary

A public key for *repladmin* defined on each node of your primary cluster must be known to *repluser* on each node of your standby cluster.

To make the key known, the directory `~repluser/.ssh` must exist on each standby node. If this directory does not exist, then create it with access only for *repluser*. Ensure that an `ls` command for the `.ssh` directory displays output similar to:

```
repluser@standby $ ls -ld ~/.ssh
drwx-----+ 1 repluser Domain Users 4096 2016-02-23 11:27 .ssh
```

If a public key for *repladmin* is defined on a given primary node, then it resides in a `.pub` file, such as `~repladmin/.ssh/id_rsa.pub`. If a public key file exists, then add its contents to the set of keys authorized to log in as *repluser* on each node of the standby where replication is run. Append the key to the file `~repluser/.ssh/authorized_keys2` on each standby node, creating this file if necessary.

If a public key file does not exist, generate a public and private key pair on the primary by running the following command as *repladmin*.

```
# ssh-keygen -t rsa
```

You can press the enter key in response to each prompt issued by the command. Copy the resulting `.pub` file to each standby node.

You have the option to share the same public/private key pair for *repladmin* across all of the nodes in your primary cluster, or to establish a different key pair for each primary node. If the same public key is valid for *repladmin* across all nodes in your primary cluster, then only that key must be added to the file `~repluser/.ssh/authorized_keys2` on each node of your standby cluster. If each primary node has its own public key for *repladmin*, then all the public keys must be added to the file. In either case, you can minimize work by copying the updated `authorized_keys2` file on a given node of the standby to the other nodes of the cluster.

Getting the standby host key

A host key for each standby node where replication may run must be known on each primary node where replication may run. One way to generate the correct key is to run `ssh` manually as the *repladmin* user from each primary node to each standby node. If the correct host key is not known already, then a warning displays and you can enable `ssh` to add the key.

Note that there are two users involved in the `ssh` connection. While `ssh` on the primary node connects to the standby node as *repladmin*, `ssh` logs in on the standby node as *repluser*. Any command run by `ssh` on the standby runs with the privileges of *repluser*, not with *repladmin* privileges.

Because the primary node connects to the standby node as user *repladmin*, the host key for the standby node must be added to the `known_hosts` file of the *repladmin* user, not the file for *repluser*. The following is an example of obtaining a host key:

```
[repladmin@primary usm]# ssh repluser@standby date
The authenticity of host 'standby (10.137.13.85)' can't be established.
RSA key fingerprint is 1b:a9:c6:68:47:b4:ec:7c:df:3a:f0:2a:6f:cf:a7:0a.
Are you sure you want to continue connecting (yes/no)?
```

If you respond with *yes*, then the `ssh` setup is complete. A host key for host standby is stored in the `known_hosts` file (`~repladmin/.ssh/known_hosts`) on the host *primary* for the user *repladmin*.

After the host key setup for standby nodes is complete on a given primary node, you need to perform an additional step if you use a Virtual IP address (VIP) to communicate with your standby cluster. You must add the VIP name or address at the start of each line of the `known_hosts` file that refers to a host in the standby cluster. For example, if you use a VIP with the name `standby12_vip`, and your `known_hosts` file contains the following two lines that refer to your standby:

```
standby1,10.242.20.22 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQAC3pM2YTd4UUUEWEoCKDGgaTgsmPkQToDrdtU+JtVIq/96mivU
BaJUK83aqzeNIQkh+hUULsUdgKoKt5bxrWYqhY6AlTEqNgBHjBrJt9C73BbQd9y48jsc2G+WQWyuI/
+s1Q+hIJdBNMxvMBQAFisPWWUcaIx9Y/Jz1PgF61RP2cbfqAzixDot9fqRrAKL3G6A75A/6TbwmEW07d1zqOv
17ZGyeDYf5zQ72F/V0P9UGMet/5DmcYTn3kTVGj0TbnrBe4A41Y4rVw5c+nZBDFre66XtORfQgwQB5ztw/Pi
08GYbcIszKoZx2HST9AZxYIagcrnNYG2Ae0K6QLxxxScP
standby2,10.242.20.23 ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQADIszcjzNtKN03SY8K1846skFTVP1HF/ykswbmktEjL6KTWTW+NR
U4MGbvkBqqdXxuPCR7aoGO2U3PEOglUVf3DWUoux8IRvqKU+dJcdTibMFkDAIhTnzbl4gz/LRTjn+GYsuP5
Qz2vgL/U0ki887mZCRjWVL1b5FNH8sXBUV2QcD7bjF98VXF6n4gd5UiIC3jv612nVTKDwtNHpUTS1dQAI+1D
tr0AieZTsuxXMaDdUZHgKDotjciMB3mCkKm/u3IFoioDqdZE4+vITX9G7DBN4CVPXawp+b5Kg8X9P+08Eehu
tM1BJ5lafy1bxoVlXUDLVIIFBJNkrsgBvxxxps7
```

To enable the use of the VIP, you would modify these two lines to read as follows:

standby12_vip,standby1,10.242.20.22 ssh-rsa

```
AAAAB3NzaC1yc2EAAAADAQABAAQ3pM2YTd4UUUEWEoCKDGgaTgsmPkQToDrdtU+JtVIq/96muivU
BaJUK83aqzeNIQkh+hUULsUdgKoKT5bXrWYqhY6AlTEqNgBHjBrJt9C73BbQd9y48jSc2G+WQWyuI/
+s1Q+hIJdBNMxvMBQafisPWWUcaIx9Y/Jz1PgF61RP2cbfqAzixDot9fqRrAKL3G6A75A/6TbwmEW07d1zqOv
l7ZGyeDYf5zQ72F/V0P9UgMEt/5DmcYTn3kTVGjOTbnRBe4A41Y4rVw5c+nZBDFre66XtOrfQgwQB5ztW/Pi
08GYbcIszKoZx2HST9AZxYIAGcrnNYG2Ae0K6QLxxxScP
```

standby12_vip,standby2,10.242.20.23 ssh-rsa

```
AAAAB3NzaC1yc2EAAAADAQABAAQDIsczjzNtKN03SY8Kl846skFTVPlHF/ykswbmktEjL6KTWTW+NR
U4MGbvKbqqdXxuPCR7aoGO2U3PEOglUVf3DWUoux8IRvqKU+dJcdTibMFkDAIhTnzb14gZ/lRTjn+GYsuP5
Qz2vgL/U0ki887mZCRjWVl1b5FNH8sXBUV2QcD7bjF98VXF6n4gd5UiIC3jv612nVTkDwtNHpUTS1dQAI+1D
tr0AieZTsuxXMaDdUZHgKDotjciMB3mCkKm/u3IFoioDqdZE4+vITX9G7DBN4CVPXawp+b5Kg8X9P+08Eehu
tMLBJ5lafy1bxoVlXUDLVIIFBJNKrsqBvxxxps7
```

Ultimately, the host key configuration performed on this first node of your primary cluster must be performed on every node in your primary cluster; the result of the above sequence, or an equivalent, must exist on each primary node. One way to minimize the manual effort required to achieve this configuration is to update the `known_hosts` file on one node of the primary cluster, then copy the updated file to the other nodes of the cluster.

Note:

By default, replication enables strict host key checking by `ssh`, to ensure that the primary node connects to the intended standby node or cluster when it runs `ssh`. However, if you are certain that this checking is unneeded, such as the case when the primary and standby clusters communicate over a private network, the use of strict host key checking by `ssh` can be disabled. For information about disabling strict host key checking, refer to the `/o sshStrictKey=no` option of the `acfsutil repl init primary` command. If strict host key checking is disabled, then no host key setup is required. For information about the `acfsutil repl init` command, refer to [acfsutil repl init](#) (page 16-69).

Adding keys for the SYSTEM user

Oracle ACFS replication uses a daemon running on a node of the primary cluster to control the progress of replication. This daemon is automatically started using the identity `NT_AUTHORITY\SYSTEM`. With this daemon, replication, using `ssh`, is able to transfer and apply primary data to the standby site.

The processes started by the daemon have the identity of the daemon. If the user is `NT_AUTHORITY\SYSTEM`, that is the user actually contacting the standby through `ssh`. `ssh` searches in the home directory for the `SYSTEM` user for a public key to present to the standby, and checks there to verify the host key presented by the standby.

To enable the daemon's successful use of `ssh`, the contents of the `.ssh` directory set up for `repladmin` on the primary should be copied to a `.ssh` directory in the home directory for `SYSTEM`. Use a location for this directory that is parallel to the locations used for other home directories.

For instance, if the home directory for user `repladmin` is `/home/repladmin`, then use `/home/SYSTEM` as the home directory for `SYSTEM`. Create the `.ssh` directory like this, while logged in as `repladmin`.

First, create the home directory for `SYSTEM` if it does not already exist:

```
$ mkdir /home/SYSTEM
```

Next, copy the `.ssh` directory for *repladmin* into the directory for *SYSTEM*:

```
$ cd $HOME
$ cp -p -R .ssh /home/SYSTEM
```

Testing Your Cygwin Setup

Now test your Cygwin setup. Using a different computer that has the `ssh` client available, execute the following command as the user *repladmin*:

```
ssh -l repluser host-address date
```

where `host-address` refers to the host where you just configured `ssh`. For example:

```
ssh -l repluser standby1.us.example.com date
```

If you have completed the installation and configuring steps successfully, the previous command should run without prompting for a password.

If you experience a process fork failure, memory leak error, or a file access error after configuring `ssh`, view the following link for a workaround:

<http://cygwin.com/faq.html>

If you are unable to find a workaround for your problem, report your problem to the Cygwin community using the following link:

<http://cygwin.com/problems.html>

18.3 Upgrading to Oracle ACFS Snapshot-Based Replication

This section describes the upgrade to Oracle ACFS snapshot-based replication. With Oracle Grid Infrastructure 12c release 2 (12.2), Oracle ACFS snapshot-based replication becomes the supported Oracle ACFS replication implementation. If you have a replication environment from a release before 12.2, then you must upgrade your existing replication installation to snapshot-based replication when you upgrade to Oracle Grid Infrastructure 12c release 2 (12.2).

Note:

- After the replication upgrade process has started, it must be completed. The process cannot be rolled back.
 - Oracle Flex ASM should be enabled, and if necessary a password file conversion performed, before upgrading to snapshot-based replication.
-
-

Oracle Flex ASM must be enabled for any existing instance when the instance is upgraded to Oracle Grid Infrastructure 12c release 2 (12.2). If a password file exists in the instance to support replication, the file must be enabled for use in the new Oracle Flex ASM instance. For example, if an instance had a previous `ORACLE_SID` set to `ASM` and a current `ORACLE_SID` set to `APX1`, then the following steps must be performed:

1. Copy the existing `orapw+ASM` file to an `+APX1` version

```
$ cp /oraroot/app/12.2.0/has0626/dbs/orapw+ASM /oraroot/app/12.2.0/has0626/dbs/
orapw+APX1
```

2. Create a new text file called `init+APX1.ora` in the `db`s directory and add only the following line to the new file:

```
remote_login_passwordfile = exclusive
```

3. Restart the Oracle ASM proxy (APX) instance.

```
$ asmcmd shutdown --target APX --immediate
```

Ensure that the environment variable `ORACLE_SID` is set to the Oracle ASM proxy (APX) instance. Then restart the instance.

```
$ asmcmd startup --pfile init+APX1.ora
```

Snapshot-based replication uses `ssh` as the transport between the primary and standby clusters. Before starting the upgrade process, you must configure `ssh` as described in [Configuring ssh for Use With Oracle ACFS Replication](#) (page 18-1).

Before upgrading to snapshot-based replication, you must upgrade your standby cluster system to Oracle Grid Infrastructure 12c release 2 (12.2). This cluster upgrade should be completed within 24 hours. The standby cluster continues to run existing replication during this upgrade process. Next, upgrade to Oracle Grid Infrastructure 12c release 2 (12.2) on the primary cluster. This upgrade should also be completed within 24 hours. During this upgrade, the primary cluster continues to run existing replication that was installed before 12.2. After the completion of the Oracle Grid Infrastructure 12c release 2 (12.2) upgrade of the standby and primary clusters, the existing replication implementation installed before 12.2 is not supported.

Immediately following the upgrade of the two clusters, you must set the `COMPATIBLE.ADVM` disk group attribute associated with the file systems to be involved in replication to `12.2.0.0.0`. Now you are ready to transition to snapshot-based replication. The transition can be made with one of the following options:

- Terminate your existing replication and initialize to use snapshot-based replication. This option is recommended if the amount of data being replicated is small (less than a few hundred gigabytes). This option may also be preferable if your primary file system cannot easily be restricted to being mounted on only one node, or if activity on the primary cannot easily be quiesced.
- Run the `acfsutil repl upgrade` command to upgrade your existing replication environment without terminating and initializing replication. This option avoids the need to transfer the entire contents of your primary file system to the standby by enabling snapshot-based replication to continue from the point where existing replication has stopped. This option is advantageous when larger amounts of data are being replicated. Because very little data is transferred by the upgrade process, you can expect the process to complete quickly.

The following list provides details on running the `acfsutil repl upgrade` command options to upgrade to snapshot-based replication.

- Before starting the replication upgrade process, you must ensure that your primary file system is mounted on only one node of the primary cluster. In addition, it is strongly recommended that you quiesce all application updates to the primary file system, then run `acfsutil repl sync apply` one final time to ensure all changes on the primary are replicated.
- The initial step in the upgrade process is to run the `acfsutil repl upgrade prepare` command on the primary cluster. This command specifies the user and host or interface name that will be used for snapshot-based replication, as well as the primary mount point. The user and host names are given with the `-s` option,

exactly as they are for the `acfsutil repl init primary` command used for snapshot-based replication.

- The next step in the process is to upgrade replication on the standby cluster, by running the `acfsutil repl upgrade standby` command. This command specifies the user to be used for snapshot-based replication, as well as the standby mount point. The user name is given with the `-u` option, exactly as it is for the `acfsutil repl init standby` command used for snapshot-based replication.
- The final step in the process is to run the `acfsutil repl upgrade primary` command on the primary cluster. This is the command that automatically terminates the previous replication deployment and initiates snapshot-based replication. This command accepts any of the command-line options accepted by the `acfsutil repl init primary` command for snapshot-based replication, except for the `-m` and `-s` options, as that information is obtained from the now-terminated previous replication environment.
- After the `acfsutil repl upgrade primary` command completes, snapshot-based replication should be active between the affected primary and standby clusters, exactly as though the `acfsutil repl init` commands had been run with snapshot-based replication in the 12.2 release. You can use the command `acfsutil repl info -c` on both the primary and standby cluster to confirm the status of replication on each cluster.

Note:

If an error occurs in running either `acfsutil repl upgrade standby` or `acfsutil repl upgrade primary`, then consider the following:

- Depending on the error, replication may appear to be uninitialized (`acfsutil repl info` may indicate that) until the upgrade command has completed successfully.
 - To continue with the upgrade, you should correct whatever error was indicated by the failing upgrade command, then simply re-issue the failing command. The command may be re-issued multiple times.
-
-

Managing Oracle Flex ASM

Oracle Flex ASM enables Oracle ASM instances to run on a separate physical server from the database servers.

This section discusses Oracle Flex ASM in the following topics:

- [Overview of Oracle Flex ASM](#) (page 19-1)
- [About Setting Up Oracle Flex ASM](#) (page 19-6)
- [Administering Oracle Flex ASM](#) (page 19-6)

See Also:

- *Oracle Clusterware Administration and Deployment Guide* for information about Oracle Flex Cluster support
 - Articles at [My Oracle Support \(https://support.oracle.com\)](https://support.oracle.com) for information about configuring Oracle Flex ASM with Oracle ASM and Oracle ACFS
 - [ASMCMD Client Cluster Management Commands](#) (page 10-30) for information about ASMCMD commands to manage client clusters
 - *Oracle Grid Infrastructure Installation Guide* for information about installing and upgrading Oracle Grid Infrastructure
-

19.1 Overview of Oracle Flex ASM

Oracle Flex ASM enables an Oracle ASM instance to run on a separate physical server from the database servers. With this deployment, larger clusters of Oracle ASM instances can support more database clients while reducing the Oracle ASM footprint for the overall system.

When using Oracle Flex ASM, Oracle ASM clients are configured with direct access to storage.

With Oracle Flex ASM, you can consolidate all the storage requirements into a single set of disk groups. All these disk groups are mounted by and managed by a small set of Oracle ASM instances running in a single cluster. You can specify the number of Oracle ASM instances with a cardinality setting. The default is three instances.

A cluster is a set of nodes that provide group membership services. Each cluster has a name that is globally unique. Every cluster has one or more Hub nodes. The Hub nodes have access to Oracle ASM disks. Every cluster has at least one private network and one public network. If the cluster is going to use Oracle ASM for storage, it has at least one Oracle ASM network. A single network can be used as both a private and an

Oracle ASM network. For security reasons, an Oracle ASM network should never be public. There can be only one Oracle Flex ASM configuration running within a cluster.

An Oracle ASM instance can operate in several configurations in Oracle Flex ASM:


- Local Oracle ASM clients with direct access to Oracle ASM disks (Standard Oracle ASM cluster)
- Oracle Flex ASM clients with direct access to Oracle ASM disks
- Oracle ACFS access through the Oracle ASM proxy instance
- Network-based connectivity to Oracle ASM disk groups with Oracle IO Server (IOS)

These configurations are illustrated in [Figure 19-1](#) (page 19-2), [Figure 19-2](#) (page 19-4), and [Figure 19-3](#) (page 19-5).

Local Oracle ASM clients with direct access to Oracle ASM disks (Standard Oracle ASM cluster)

With this mode, illustrated as **Hub Node A** in [Figure 19-1](#) (page 19-2), Oracle ASM continues to support existing standard architecture in which database clients are running with an Oracle ASM instance on the same host computer. The local client architecture is only supported on a Hub node.

Figure 19-1 Oracle Flex ASM Client Configurations



In this configuration, the database instances are on the same Hub node as the Oracle ASM instance and are referred to as local Oracle ASM client instances. Oracle ASM

metadata moves between Oracle ASM and the database instances. This client has direct I/O access to Oracle ASM disks.

Local mode does not use Oracle Flex ASM, so clusters configured with local Oracle ASM do not require an Oracle ASM network, nor do they contain other Oracle Flex ASM services.

Oracle Flex ASM clients with direct access to Oracle ASM disks

With this mode, illustrated as **Hub Node B** in [Figure 19-1](#) (page 19-2), database clients that are running on Hub nodes of the Oracle ASM cluster access Oracle ASM remotely for metadata, but perform block I/O operations directly to Oracle ASM disks. The hosts running the Oracle ASM server and the remote database client must both be Hub nodes. A Hub node is a node in an Oracle ASM cluster that is tightly connected with other servers and has direct access to a shared disk.

In this configuration, the database instances are on different host computers than the nearby Oracle ASM instance, as shown on **Hub Node C** in [Figure 19-1](#) (page 19-2), and are referred to as Oracle ASM client instances.

The databases are in the same Oracle ASM cluster as the Oracle ASM instance and the database instances are located on a Hub node. Oracle ASM metadata moves between Oracle ASM and the database instance. This client has direct I/O access to Oracle ASM disks.


Depending on the distribution of database instances and Oracle ASM instances, a database client may access Oracle ASM locally on the same node or remotely over the Oracle ASM network. This mode of operation is used by database clients on Hub nodes in the Oracle ASM cluster. Direct access mode is also the only Oracle Flex ASM configuration supported by Oracle ASM cluster file system.

Oracle ACFS access through the Oracle ASM proxy instance

An Oracle ASM proxy instance is an Oracle instance running on a Hub node with a direct Oracle ASM client. An Oracle ASM proxy instance provides support for Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM).

[Figure 19-2](#) (page 19-4) shows the configuration of Oracle ACFS and Oracle ADVM with an Oracle ASM Proxy server in Oracle Flex ASM.

Figure 19-2 Oracle ACFS and Oracle ADVM in Oracle Flex ASM Configuration

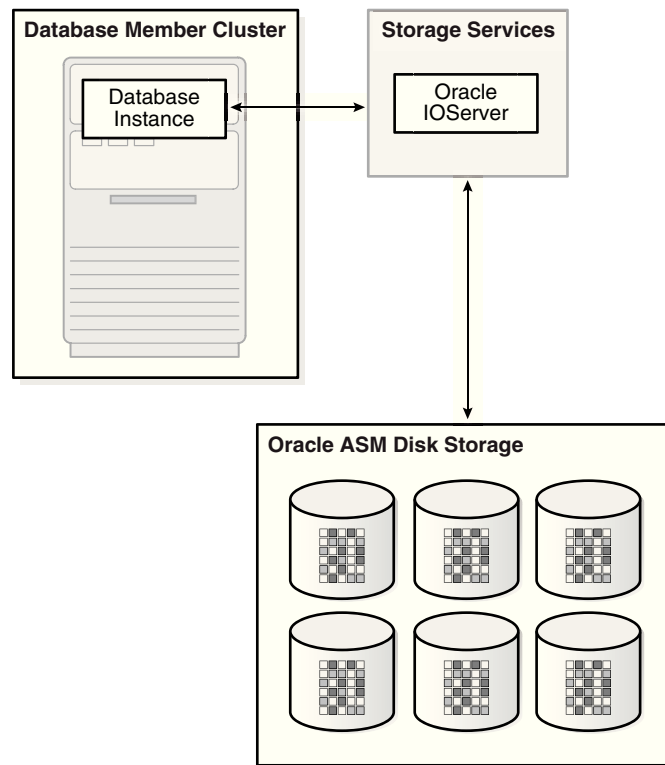


The `INSTANCE_TYPE` initialization parameter is set to `ASMPROXY` for Oracle ASM proxy instances.

Network-based connectivity to Oracle ASM disk groups with Oracle IO Server (IOS)

An Oracle IO Server instance provides Oracle ASM file access for Oracle Database instances on nodes of Oracle member clusters that do not have connectivity to Oracle ASM managed disks.

Figure 19-3 (page 19-5) shows the relationship of Oracle IO Server and an Oracle Database in a member cluster.

Figure 19-3 Oracle IO Server Configuration

In addition, IOS enables you to configure client clusters on such nodes. On the storage cluster, an IO Server instance opens up network ports to which clients send their IO. The IO Server instance receives data packets from the client and performs the appropriate IO to Oracle ASM disks similar to any other database client. On the client side, databases can use dNFS to communicate with an IO Server instance. However, there is no client side configuration so you are not required to provide a server IP address or any additional configuration information. On nodes and clusters that are configured to access Oracle ASM files through IO Server, the discovery of the Oracle IOS instance occurs automatically.

The IOS instance contains the following processes:

- **Network processes:** These processes open up network ports and receive dNFS requests from the client side. These requests are then placed in a queue for IO and Identifier processes to pick up and process. Network processes also pick up responses from those other types of processes and send the response back to the client.
- **Identifier processes:** These processes pick up dNFS requests to create, delete, identify, un-identify, and resize Oracle ASM files.
- **IO processes:** These processes perform the actual IO on files that are identified by the local IOS instance.

The `ASM_IO_PROCESSES` initialization parameter specifies the number of IO worker processes to be started in an Oracle IO server instance. For information about the `ASM_IO_PROCESSES` initialization parameter, refer to [ASM_IO_PROCESSES](#) (page 3-11).

19.2 About Setting Up Oracle Flex ASM

To install an Oracle Flex ASM deployment, categorize the networks and choose the list of networks for use as Oracle ASM networks.

If you choose Oracle Flex ASM during a new installation, OUI requires you to choose the Oracle ASM networks.

The Oracle ASM listener resource is automatically created for each Oracle ASM network and then started on all nodes.

See Also:

Oracle Grid Infrastructure Installation Guide for information about Oracle Clusterware installation

19.3 Administering Oracle Flex ASM

Oracle Flex ASM is managed by ASMCA, CRSCTL, SQL*Plus, and SRVCTL. The `INSTANCE_TYPE` initialization parameter specifies the type of instance.

The `INSTANCE_TYPE` initialization parameter has an additional value `ASMPROXY`, in addition to `ASM` and `RDBMS`, to identify Oracle ASM proxy instances. An Oracle ASM proxy instance has its parameter set to `ASMPROXY`.

You can use the `ASMCMD showclustermode` command to determine whether Oracle Flex ASM is enabled. For example:

```
$ asmcmd showclustermode
ASM cluster : Flex mode enabled
```

`SRVCTL` is extended to enable an administrator to create or change attributes of Oracle Clusterware resources. You can use `SRVCTL` to determine the status of the instances in an Oracle Flex ASM configuration. For example:

```
$ srvctl status asm -detail
ASM is running on mynoden02,mynoden01
ASM is enabled.
```

You can also use `SRVCTL` to determine whether Oracle Flex ASM is enabled. If enabled, then `srvctl config asm` displays the number of Oracle ASM instances that has been specified for use with the Oracle Flex ASM configuration. For example:

```
$ srvctl config asm
ASM instance count: 3
```

You can modify the Oracle ASM instance count, or cardinality, with the `SRVCTL modify asm` command. For example:

```
$ srvctl modify asm -count 4

$ srvctl modify asm -count ALL
```

You can view Oracle Flex ASM connections with `SQL*Plus` and `ASMCMD` commands. For example:

```
SQL> SELECT instance_name, db_name, status FROM V$ASM_CLIENT;
INSTANCE_NAME  DB_NAME  STATUS
```

```

-----
+ASM1          +ASM    CONNECTED
orcl1          orcl    CONNECTED
orcl2          orcl    CONNECTED

$ asmcmd lsct data
DB_Name Status      Software_Version  Compatible_version  Instance_Name  Disk_Group
+ASM     CONNECTED      12.1.0.0.2        12.1.0.0.2         +ASM           DATA
orcl     CONNECTED      12.1.0.0.2        12.0.0.0.0         orcl1          DATA
orcl     CONNECTED      12.1.0.0.2        12.0.0.0.0         orcl2          DATA

```

Clients are automatically relocated to another instance if an Oracle ASM instance fails. If necessary, clients can be manually relocated.

The `SRVCTL UPDATE INSTANCE` and `SRVCTL UPDATE IOSERVER` commands can change the Oracle ASM instance for a database, or the Oracle IO Server instance for a database, or the Oracle ASM instance for an Oracle IO Server. For example:

```

$ srvctl update instance -db my_orcl_db -instance my_instance_1 -targetinstance
my_asm_instance_1

$ srvctl update ioserver -instance my_instance_2 -targetinstance my_asm_instance_2

```

You can run the `SRVCTL STATUS` commands with the `-detail` option to display which Oracle ASM instance or Oracle IO server instance each client is connected to.

You can also use the `ALTER SYSTEM RELOCATE CLIENT` command to relocate a client. For example:

```
SQL> ALTER SYSTEM RELOCATE CLIENT 'client-id';
```

In the previous SQL statement, *client-id* is of the form *instance_name:db_name*. The `INSTANCE_NAME` and `DB_NAME` columns are contained in the `V$ASM_CLIENT` view. You must connect as `SYSASM` to the Oracle ASM instance to run the SQL statement. When you issue this statement, the connection to the client is terminated and the client fails over to the least loaded instance. If the client is currently connected to the least loaded instance, then the connection to the client is terminated and the client fails over to that same instance.

Every database user must have a wallet with credentials to connect to Oracle ASM. `CRSCTL` commands can be used by the database user to manage this wallet. All Oracle ASM user names and passwords are system generated.

There are no new initialization parameters specifically for instances in an Oracle Flex ASM configuration; however, the settings of existing parameters should be reviewed and possibly adjusted for the Oracle Flex ASM environment. Refer to "[Recommended Settings for Oracle ASM Initialization Parameters](#) (page 3-8)".

See Also:

- *Oracle Database Reference* for more information about the `INSTANCE_TYPE` initialization parameter
 - *Oracle Database SQL Language Reference* for information about the `ALTER SYSTEM` statement
 - *Oracle Clusterware Administration and Deployment Guide* for information about `CRSCTL` commands
 - *Oracle Real Application Clusters Administration and Deployment Guide* for information about `SRVCTL` commands
 - *Oracle Grid Infrastructure Installation Guide* for information about installing Oracle Clusterware
-
-

Administering Oracle ASM Filter Driver

Oracle ASM Filter Driver (Oracle ASMFD) simplifies the configuration and management of disk devices by eliminating the need to rebind disk devices used with Oracle ASM each time the system is restarted.

This section contains the following topics:

- [About Oracle ASM Filter Driver](#) (page 20-2)
- [About Oracle ASMLIB and Oracle ASM Filter Driver](#) (page 20-3)
- [About Partition Tables and Oracle ASM Filter Driver](#) (page 20-3)
- [Configuring Oracle ASM Filter Driver](#) (page 20-3)
- [About Migrating to Oracle ASM Filter Driver From ASMLIB](#) (page 20-7)
- [Determining Whether Oracle ASM Filter Driver Has Been Configured](#) (page 20-10)
- [Updating the Oracle ASM Filter Driver AFD_DISKSTRING Parameter](#) (page 20-10)
- [Updating the Oracle ASM ASM_DISKSTRING Parameter for Oracle ASM Filter Driver Disks](#) (page 20-11)
- [Setting, Clearing, and Scanning Oracle ASM Filter Driver Labels](#) (page 20-11)
- [Deconfiguring Oracle ASM Filter Driver](#) (page 20-12)
- [Oracle ASM Filter Driver Commands](#) (page 20-15)

Note:

- The configuration for Oracle ASMFD can now be enabled with a check box to be an automated process during Oracle Grid Infrastructure installation.
 - If ASMLIB is installed and configured for an existing Oracle ASM installation, then you must explicitly uninstall the existing ASMLIB configuration before installing and configuring Oracle ASMFD.
 - Oracle ASM Filter Driver (Oracle ASMFD) is available on Linux systems starting with Oracle Database 12c Release 1 (12.1.0.2). Oracle ASMFD is available on Solaris systems starting with Oracle Database 12c Release 2 (12.2.0.1).
 - In the steps of the procedures described in the topics of this section, the `$ORACLE_HOME` environmental variable is set to the directory path of the Oracle Grid Infrastructure home. Commands that show # as the operating system prompt must be run as the root user. Commands that show \$ as the operating system prompt should be run as the owner of Oracle Grid Infrastructure home
-
-

See Also:

- [ASMCMD Oracle ASM Filter Driver Management Commands](#) (page 10-32) for information about the ASMCMD commands for administering Oracle ASMFD
 - [About ASMCMD](#) (page 10-1) for information about all the ASMCMD commands
 - [Managing Oracle ASM Filter Driver With Oracle Enterprise Manager](#) (page 7-20) for information about using Oracle Enterprise Manager to administer Oracle ASMFD
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for your operating system for information about installing and configuring Oracle Grid Infrastructure
 - *Oracle Grid Infrastructure Installation and Upgrade Guide* for your operating system for information about deinstalling Oracle ASMLIB
 - *Oracle Database 2 Day + Real Application Clusters Guide* for information about installing and configuring Oracle Grid Infrastructure
-
-

20.1 About Oracle ASM Filter Driver

Oracle ASM Filter Driver (Oracle ASMFD) is a kernel module that resides in the I/O path of the Oracle ASM disks. Oracle ASM uses the filter driver to validate write I/O requests to Oracle ASM disks.

Oracle ASMFD simplifies the configuration and management of disk devices by eliminating the need to rebind disk devices used with Oracle ASM each time the system is restarted.

Oracle ASM Filter Driver rejects any I/O requests that are invalid. This action eliminates accidental overwrites of Oracle ASM disks that would cause corruption in the disks and files within the disk group. For example, the Oracle ASM Filter Driver filters out all non-Oracle I/Os which could cause accidental overwrites.

20.2 About Oracle ASMLIB and Oracle ASM Filter Driver

Starting with Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle ASM Filter Driver (Oracle ASMFd) cannot be installed and configured if Oracle ASMLIB has been installed on the system. If you want to install and configure Oracle ASMFd, you must first uninstall Oracle ASMLIB.

See Also:

Oracle Grid Infrastructure Installation Guide for your operating system for information about deinstalling Oracle ASMLIB

20.3 About Partition Tables and Oracle ASM Filter Driver

Extended partition tables are not supported with Oracle ASM filter driver (ASMFd) in Oracle Automatic Storage Management 12.2 Release 1 (12.2.0.1).

20.4 Configuring Oracle ASM Filter Driver

You can configure Oracle ASM Filter Driver (Oracle ASMFd) for device persistence during the installation of Oracle Grid Infrastructure or after Oracle Grid Infrastructure has been installed.

- [Configuring Oracle ASM Filter Driver During Installation](#) (page 20-4)
- [Configuring Oracle ASM Filter Driver After Installation](#) (page 20-5)

Note:

- If ASMLIB is installed and configured for an existing Oracle ASM installation, then you must explicitly deinstall the existing ASMLIB configuration before installing and configuring Oracle ASMFd. Refer to *Oracle Grid Infrastructure Installation Guide* for your operating system for information about deinstalling Oracle ASMLIB and the installation of Oracle Grid Infrastructure.
 - If ASMLIB has been previously installed on the system, then you must set the Oracle ASM disk string to '' (with ASMCMD `dsset ''`) before running ASMCMD `afd_configure` to configure Oracle ASMFd.
-
-

See Also:

- *Oracle Grid Infrastructure Installation Guide* for your operating system for information about installing and configuring Oracle Grid Infrastructure
 - *Oracle Oracle Database 2 Day + Real Application Clusters Guide* for information about installing and configuring Oracle Grid Infrastructure
-
-

20.4.1 Configuring Oracle ASM Filter Driver During Installation

During the installation of Oracle Grid Infrastructure, you can optionally enable automated installation and configuration of Oracle ASM Filter Driver for your system with the **Configure ASM Filter Driver** check box on the **Create ASM Disk Group** wizard page. When you enable the **Configure ASM Filter Driver** box, an automated process for Oracle ASMFD is launched during Oracle Grid Infrastructure installation.

If you do not use udev on the system where the Oracle Grid Infrastructure is installed, then you can also complete the following procedure to provision disks for Oracle ASMFD before the installer is launched. This procedure must be performed after the Oracle Grid Infrastructure installation zip files have been unzipped in the Oracle Grid Infrastructure home, but before you enable the **Configure ASM Filter Driver** check box during the installation process when the setup script is run.

1. To configure shared disks for use with Oracle ASM Filter Driver, log in as the root user and set the environment variable \$ORACLE_HOME to the location of the Grid home and the environment variable \$ORACLE_BASE to a temporary location

```
su root
# set ORACLE_HOME=/u01/app/oracle/12.2.0/grid
# set ORACLE_BASE=/tmp
```

The ORACLE_BASE variable is set to a temporary location to avoid creating diagnostic or trace files in the Grid home before the Oracle Grid Infrastructure installation.

Before running the commands in the following steps, ensure that you are in the \$ORACLE_HOME/bin directory.

2. Use the ASMCMD afd_label command to provision disk devices for use with Oracle ASM Filter Driver as follows.

```
# asmcmd afd_label DATA1 /dev/diskla --init
```

3. Use the ASMCMD afd_ls1b1 command to verify the device has been marked for use with Oracle ASMFD.

```
# asmcmd afd_ls1b1 /dev/diskla
```

4. Unset the ORACLE_BASE variable when you have finished provisioning disks for Oracle ASMFD.

```
# unset ORACLE_BASE
```

5. Run the setup installation script (gridSetup.sh) to install Oracle Grid Infrastructure and enable Oracle ASM Filter Driver configuration.

On the **Create ASM Disk Group** wizard page during the installation process, check the **Configure ASM Filter Driver** box to enable configuration of Oracle ASM Filter Driver and then add your provisioned disks to the initial disk group (typically named DATA).

For information about the ASMCMD commands for administering Oracle ASMFD, refer to "[ASMCMD Oracle ASM Filter Driver Management Commands](#) (page 10-32)".

20.4.2 Configuring Oracle ASM Filter Driver After Installation

If you did not enable configuration of Oracle ASMFD during installation, you can manually configure your Oracle ASM devices to use Oracle ASMFD, as described in the following topics.

- ["Configuring Oracle ASM in an Oracle Grid Infrastructure Clusterware Environment \(page 20-5\)"](#)
- ["Configuring Oracle ASM in an Oracle Grid Infrastructure Standalone \(Oracle Restart\) Environment \(page 20-6\)"](#)

See Also:

Oracle Clusterware Administration and Deployment Guide for information about using CRSCTL commands

20.4.2.1 Configuring Oracle ASM in an Oracle Grid Infrastructure Clusterware Environment

To configure Oracle ASMFD in an Oracle Clusterware environment, follow these steps:

1. As the Oracle Grid Infrastructure owner update the Oracle ASM disk discovery string to enable Oracle ASMFD to discover devices in the future.

For example, check the current value of the Oracle ASM disk discovery string and then update the value.

```
$ $ORACLE_HOME/bin/asmcmd dsget
```

```
$ $ORACLE_HOME/bin/asmcmd dsset old_diskstring, 'AFD:*
```

The value of *old_diskstring* is the current Oracle ASM disk discovery string value.

2. As the Oracle Grid Infrastructure owner list the nodes and node roles in your cluster:

```
$ $ORACLE_HOME/bin/olsnodes -a
```

3. On each Hub and Leaf node, do the following, either in rolling or non-rolling mode:

- a. Log in as the `root` user and stop Oracle Grid Infrastructure:

```
# $ORACLE_HOME/bin/crsctl stop crs
```

If the command returns an error, then stop Oracle Grid Infrastructure forcibly as follows:

```
# $ORACLE_HOME/bin/crsctl stop crs -f
```

- b. As `root`, configure Oracle ASMFD to filter at the node level:

```
# $ORACLE_HOME/bin/asmcmd afd_configure
```

- c. As the Oracle Grid Infrastructure owner verify the status of Oracle ASMFD:

```
$ $ORACLE_HOME/bin/asmcmd afd_state
ASMCMD-9526: The AFD state is 'LOADED' and filtering is 'ENABLED' on host
'myhost'
```

- d. As root, start the Oracle Clusterware stack on the node:

```
# $ORACLE_HOME/bin/crsctl start crs
```

- e. As the Oracle Grid Infrastructure owner set the Oracle ASMFD discovery disk string to the original Oracle ASM disk discovery string value that was retrieved in Step 1:

```
$ $ORACLE_HOME/bin/asmcmd afd_dsset old_diskstring
```

The value of *old_diskstring* is the old disk discovery string value without the AFD: (Oracle ASMFD) path.

Oracle ASM Filter Driver should identify and start managing disks previously managed by Oracle ASMLIB.

See Also:

- [Updating the Oracle ASM ASM_DISKSTRING Parameter for Oracle ASM Filter Driver Disks](#) (page 20-11) for information about updating the Oracle ASM discovery string
 - [Determining Whether Oracle ASM Filter Driver Has Been Configured](#) (page 20-10) for information about checking on the state of the Oracle ASM Filter Driver
 - [Updating the Oracle ASM Filter Driver AFD_DISKSTRING Parameter](#) (page 20-10) for information about updating the Oracle ASM Filter Driver discovery disk discovery string
-
-

20.4.2.2 Configuring Oracle ASM in an Oracle Grid Infrastructure Standalone (Oracle Restart) Environment

To configure Oracle ASMFD in a standalone environment, follow these steps:

1. As the Oracle Grid Infrastructure standalone server owner update the Oracle ASM disk discovery string to enable Oracle ASMFD to discover disk devices.

For example, check the current value of the Oracle ASM disk discovery string and then update the value.

```
$ $ORACLE_HOME/bin/asmcmd dsget

$ $ORACLE_HOME/bin/asmcmd dsset old_diskstring, 'AFD:*
```

Where *old_diskstring* is the current disk discovery string value.

2. Log in as the root user and stop Oracle Grid Infrastructure for a standalone server using the following command:

```
# $ORACLE_HOME/bin/crsctl stop has -f
```

3. As root, configure Oracle ASMFD using the following command:

```
# $ORACLE_HOME/bin/asmcmd afd_configure
```

This command configures Oracle ASMFD and deconfigures Oracle ASMLIB, if it exists.

4. As the Oracle Grid Infrastructure standalone server owner verify the Oracle ASMFD status:

```
$ $ORACLE_HOME/bin/asmcmd afd_state
```

```
ASMCMD-9526: The AFD state is 'LOADED' and filtering is 'ENABLED' on host 'myhost'
```

5. As root, start Oracle Grid Infrastructure for a standalone server:

```
# $ORACLE_HOME/bin/crsctl start has
```

6. As the Oracle Grid Infrastructure standalone server owner set the Oracle ASMFD disk discovery string to the original value of the Oracle ASM disk discovery string that was retrieved in Step 1:

```
$ $ORACLE_HOME/bin/asmcmd afd_dsset disk_string
```

Oracle ASM Filter Driver should identify and start managing disks previously managed by Oracle ASMLIB.

See Also:

- [Updating the Oracle ASM ASM_DISKSTRING Parameter for Oracle ASM Filter Driver Disks](#) (page 20-11) for information about updating the Oracle ASM discovery string
 - [Determining Whether Oracle ASM Filter Driver Has Been Configured](#) (page 20-10) for information about checking on the state of the Oracle ASM Filter Driver
-
-

20.5 About Migrating to Oracle ASM Filter Driver From ASMLIB

If Oracle ASMLIB was installed, but not used earlier, you must create disk labels to enable migration of Oracle ASM disk groups to Oracle ASM Filter Driver (Oracle ASMFD) after installing Oracle Grid Infrastructure 12c Release 1 (12.1.0.2).

Oracle recommends that you temporarily move Oracle Cluster Registry (OCR) and voting files to another disk group if one is available, as described in "[Migrating Oracle ASM Disk Groups without Oracle Cluster Registry or Voting Files](#) (page 20-8)", and migrate the disk group to use Oracle ASMFD. After migrating the disk group to use Oracle ASMFD, move OCR and voting files back to the disk group. You can similarly migrate any other disk groups if they contain OCR or voting files to ensure online migration of all disk groups to Oracle ASMFD.

This section contains the following topics:

- "[Migrating Oracle ASM Disk Groups without Oracle Cluster Registry or Voting Files](#) (page 20-8)"
- "[Migrating Oracle ASM Disk Groups with Oracle Cluster Registry or Voting Files](#) (page 20-8)"
- "[Migrating Oracle ASM Disk Groups in an Oracle Grid Infrastructure Standalone \(Oracle Restart\) Environment](#) (page 20-9)"

See Also:

- *Oracle Clusterware Administration and Deployment Guide* for information about replacing Oracle Cluster Registry locations and migrating voting files
 - *Oracle Clusterware Administration and Deployment Guide* for information about using CRSCTL commands
 - *Oracle Clusterware Administration and Deployment Guide* for information about managing voting files, Oracle Cluster Registry (OCR), and Oracle Local Registry (OLR)
 - *Oracle Real Application Clusters Administration and Deployment Guide* for information about SRVCTL commands
-
-

20.5.1 Migrating Oracle ASM Disk Groups without Oracle Cluster Registry or Voting Files

To migrate Oracle ASM disk groups without Oracle Cluster Registry (OCR) or voting files to Oracle ASMFD:

1. Log in as the Oracle Grid Infrastructure owner on any node to run the commands in this procedure.

2. List the existing disk groups:

```
$ $ORACLE_HOME/bin/asmcmd lsdg
```

3. List the associated disks:

```
$ $ORACLE_HOME/bin/asmcmd lsdsk -G diskgroup
```

4. Check if Oracle ASM is active:

```
$ $ORACLE_HOME/bin/srvctl status asm
```

5. Stop the databases and dismount the disk group on all nodes:

```
$ $ORACLE_HOME/bin/srvctl stop diskgroup -diskgroup diskgroup -f
```

6. Label all existing disks in the disk group by running the following command for each disk on a Hub node:

```
$ $ORACLE_HOME/bin/asmcmd afd_label label disk_path --migrate
```

7. Scan the disks on all Hub nodes:

```
$ $ORACLE_HOME/bin/asmcmd afd_scan
```

8. Start the databases and mount the disk group on all nodes:

```
$ $ORACLE_HOME/bin/srvctl start diskgroup -diskgroup diskgroup
```

20.5.2 Migrating Oracle ASM Disk Groups with Oracle Cluster Registry or Voting Files

To migrate Oracle ASM disk groups with Oracle Cluster Registry (OCR) or voting files to Oracle ASM Filter Driver (Oracle ASMFD):

1. Log in as the `root` user and list the disk groups with OCR and voting files by running the following commands on one node:

```
# $ORACLE_HOME/bin/ocrcheck -config
# $ORACLE_HOME/bin/crsctl query css votedisk
```

2. As the Oracle Grid Infrastructure owner list the disks associated with the disk groups:

```
$ $ORACLE_HOME/bin/asmcmd lsdsk -G disk_group
```

3. As `root`, stop the databases and Oracle Clusterware on all nodes:

```
# $ORACLE_HOME/bin/crsctl stop cluster -all
```

4. As the Oracle Grid Infrastructure owner label all existing disks in the disk group by running the following command for each disk on a Hub node:

```
$ $ORACLE_HOME/bin/asmcmd afd_label label disk_path --migrate
```

5. As the Oracle Grid Infrastructure owner rescan the disks on all Hub nodes by running the following command on all of the Hub nodes:

```
$ $ORACLE_HOME/bin/asmcmd afd_scan
```

6. As `root`, start the Oracle Clusterware stack on all nodes and mount the OCR and voting files disk groups and databases:

```
# $ORACLE_HOME/bin/crsctl start cluster -all
```

20.5.3 Migrating Oracle ASM Disk Groups in an Oracle Grid Infrastructure Standalone (Oracle Restart) Environment

To migrate your existing Oracle ASM disk groups to Oracle ASM Filter Driver in an Oracle Grid Infrastructure standalone environment, perform the following steps:

1. Log in as the Oracle Grid Infrastructure standalone server owner to run the steps in this procedure.

2. List the existing disk groups:

```
$ $ORACLE_HOME/bin/asmcmd lsdg
```

3. List the existing disks:

```
$ $ORACLE_HOME/bin/asmcmd lsdsk -G diskgroup_name
```

4. Check if the status of your Oracle ASM instance is active:

```
$ $ORACLE_HOME/bin/srvctl status asm
```

5. Stop all the databases and dismount all disk groups. For each database, run the following commands:

```
$ $ORACLE_HOME/bin/srvctl stop database -db db_unique_name
```

```
$ /$ORACLE_HOME/bin/srvctl stop diskgroup -diskgroup diskgroup_name -f
```

6. Label all existing disks in the disk group by running the following command for each disk:

```
$ $ORACLE_HOME/bin/asmcmd afd_label label disk_path --migrate
```

7. Rescan the disks:

```
$ $ORACLE_HOME/bin/asmcmd afd_scan
```

8. Start the database and mount the disk group:

```
$ $ORACLE_HOME/bin/srvctl start diskgroup -diskgroup diskgroup_name
```

20.6 Determining Whether Oracle ASM Filter Driver Has Been Configured

The value of the `AFD_STATE` parameter of `SYS_ASMFD_PROPERTIES` specifies whether Oracle ASMFD is configured for the Oracle ASM instance.

You can check the state of Oracle ASMFD with the `ASMCMD afd_state` command. For example:

```
$ $ORACLE_HOME/bin/asmcmd afd_state
ASMCMD-9526: The AFD state is 'LOADED' and filtering is 'DEFAULT' on host 'myhost'
```

For information about using the `ASMCMD afd_state` command to determine the state of Oracle ASMFD, refer to "[afd_state](#) (page 10-40)".

To determine if Oracle's ASMFD is present on the host, you can also display the value of `AFD_STATE` from `SYS_CONTEXT`. You must run the query on the Oracle ASM instance.

If the value of `AFD_STATE` is equal to `NOT AVAILABLE`, then Oracle ASMFD is not configured.

```
SQL> SELECT SYS_CONTEXT('SYS_ASMFD_PROPERTIES', 'AFD_STATE') FROM DUAL;
SYS_CONTEXT('SYS_ASMFD_PROPERTIES', 'AFD_STATE')
-----
NOT AVAILABLE
```

A value of `CONFIGURED` means that Oracle ASMFD is completely set up and the Oracle ASM instance can register with the driver.

```
SQL> SELECT SYS_CONTEXT('SYS_ASMFD_PROPERTIES', 'AFD_STATE') FROM DUAL;
SYS_CONTEXT('SYS_ASMFD_PROPERTIES', 'AFD_STATE')
-----
CONFIGURED
```

20.7 Updating the Oracle ASM Filter Driver `AFD_DISKSTRING` Parameter

The `AFD_DISKSTRING` parameter specifies the Oracle ASMFD disk discovery string that is used to identify the disks to be managed by Oracle ASMFD.

You can set and display the `AFD_DISKSTRING` parameter with the `ASMCMD afd_dsset` and `afd_dsget` commands. For example:

```
$ $ORACLE_HOME/bin/asmcmd afd_dsset '/dev/rdisk/mydisks/*'

$ $ORACLE_HOME/bin/asmcmd afd_dsget
AFD discovery string: /dev/rdisk/mydisks/*
```

For information about `ASMCMD` commands to display and set the Oracle ASMFD disk discovery string, refer to "[afd_dsget](#) (page 10-36)" and "[afd_dsset](#) (page 10-36)".

You can also set the AFD_DISKSTRING with the ALTER SYSTEM SQL statement. A label is created in the disk header of those disks identified by the Oracle ASMFD disk discovery string.

```
SQL> ALTER SYSTEM AFD_DISKSTRING SET '/dev/disk0','/dev/disk1','/devices/dsk/*';
System altered.
```

You can retrieve the value of AFD_DISKSTRING parameter with the following query.

```
SQL> SELECT SYS_CONTEXT('SYS_ASMFD_PROPERTIES', 'AFD_DISKSTRING') FROM DUAL;
SYS_CONTEXT('SYS_ASMFD_PROPERTIES', 'AFD_DISKSTRING')
```

```
-----
'/dev/disk0','/dev/disk1','/devices/dsk/*'
```

20.8 Updating the Oracle ASM ASM_DISKSTRING Parameter for Oracle ASM Filter Driver Disks

You can update the Oracle ASM disk discovery string to add or remove Oracle ASMFD disk label names to and from the ASM_DISKSTRING initialization parameter.

For example, you can add the Oracle ASMFD disks to the ASM_DISKSTRING initialization parameter as follows:

```
ASM_DISKSTRING = 'AFD:DISK0', 'AFD:DISK1', '/dev/rdisk/mydisks/*'
```

Or you can set the ASM_DISKSTRING initialization parameter as follows:

```
ASM_DISKSTRING = 'AFD:*', '/dev/rdisk/mydisks/*'
```

You can display and set the Oracle ASM disk discovery string with the ASMCMD dsget and dsset commands. For example, you can set the Oracle ASM disk discovery string to add Oracle ASMFD disks as follows:

```
$ $ORACLE_HOME/bin/asmcmd dsset 'AFD:*','/dev/rdisk/mydisks/*'
```

You can remove previously added Oracle ASMFD disks as follows:

```
$ $ORACLE_HOME/bin/asmcmd dsset '/dev/rdisk/mydisks/*'
```

For information about ASMCMD commands to display and set the Oracle ASM disk discovery string, refer to "[dsget](#) (page 10-13)" and "[dsset](#) (page 10-13)".

For information about the ASM_DISKSTRING initialization parameter, refer to "[ASM_DISKSTRING](#) (page 3-10)".

20.9 Setting, Clearing, and Scanning Oracle ASM Filter Driver Labels

Setting a label provisions a disk to be used by Oracle ASMFD. After the label is set, the specified disk is managed by Oracle ASMFD.

You can add, remove, and scan labels with the ASMCMD afd_label, afd_unlabel, and afd_scan commands. For example:

```
$ $ORACLE_HOME/bin/asmcmd afd_label 'disk0' '/dev/rdisk/mydisks/disk0'
```

```
$ $ORACLE_HOME/bin/asmcmd afd_unlabel 'disk0'
```

```
$ $ORACLE_HOME/bin/asmcmd afd_scan '/dev/rdisk/mydisks/*'
```

For information about ASMCMD commands to add and remove labels on Oracle ASMFD disks, refer to "[afd_label](#) (page 10-37)", "[afd_unlabel](#) (page 10-41)", and

"afd_scan (page 10-40)". In addition, ASMCA provides support for adding and removing labels on Oracle ASMFD disks. For information about using ASMCA to administer disk groups, refer to "[Managing Disk Groups with ASMCA \(page 9-2\)](#)".

You can also manage labels with SQL statements. You can set a label with the ALTER SYSTEM LABEL SET SQL statement. For example:

```
SQL> ALTER SYSTEM LABEL SET 'disk0' TO '/dev/disk0';
System altered.
```

```
SQL> SELECT UPPER(path) FROM V$ASM_DISK ORDER BY PATH;
UPPER(PATH)
```

```
-----
AFD:DISK0
```

When you run the statement, you can use the optional RENAME or MIGRATE option. If a disk was previously provisioned for Oracle ASMFD, you can rename the label with the RENAME option. Note that the device should not be managed with Oracle ASMFD when the command is run. If a disk was previously used for an Oracle ASM disk group and the disk group has been dismounted, then you can label this disk using the MIGRATE option.

You can use ALTER SYSTEM LABEL CLEAR to remove the label from a device and stop Oracle ASMFD from managing the device. For example:

```
SQL> ALTER SYSTEM LABEL CLEAR 'disk0';
System altered.
```

You can use ALTER SYSTEM LABEL SCAN on remote nodes after the ALTER SYSTEM LABEL SET command is run on the local node.

Because ALTER SYSTEM LABEL SET statement writes the label on the disk header and the disk is shared across nodes, the same statement is not run on other nodes of the cluster.

If the device-path is not specified, then the statement uses the AFD_DISKSTRING parameter value to perform the scan operation.

```
SQL> ALTER SYSTEM LABEL SCAN
```

20.10 Deconfiguring Oracle ASM Filter Driver

You can deconfigure Oracle ASM Filter Driver (Oracle ASMFD) if it has been configured on your system.

This section contains the following topics:

- "[Deconfiguring Oracle ASM Filter Driver in an Oracle Grid Infrastructure Clusterware Environment \(page 20-13\)](#)"
- "[Deconfiguring Oracle ASM Filter Driver in an Oracle Grid Infrastructure Standalone \(Oracle Restart\) Environment \(page 20-14\)](#)"

For information about the ASMCMD commands for administering Oracle ASMFD, refer to "[ASMCMD Oracle ASM Filter Driver Management Commands \(page 10-32\)](#)".

See Also:

Oracle Clusterware Administration and Deployment Guide for information about using CRSCTL commands

20.10.1 Deconfiguring Oracle ASM Filter Driver in an Oracle Grid Infrastructure Clusterware Environment

Perform the following steps to deconfigure Oracle ASM Filter Driver in an Oracle Clusterware environment:

1. Update the Oracle ASM disk discovery string to enable Oracle ASM to discover disk devices after Oracle ASMFD is deconfigured.
2. As the Oracle Grid Infrastructure owner list the nodes and node roles in your cluster by running the following command on any node:

```
$ $ORACLE_HOME/bin/olsnodes -a
```

3. On each Hub and Leaf node, do the following, either in rolling or non-rolling mode:

- a. Log in as the `root` user and stop Oracle Grid Infrastructure:

```
# $ORACLE_HOME/bin/crsctl stop crs
```

If the command returns any error, then stop Oracle Grid Infrastructure forcibly as follows:

```
# $ORACLE_HOME/bin/crsctl stop crs -f
```

- b. As `root`, stop Oracle ACFS kernel modules to ensure the most reliable shutdown:

```
# $ORACLE_HOME/bin/acfsload stop
```

- c. As `root`, deconfigure Oracle ASMFD:

```
# $ORACLE_HOME/bin/asmcmd afd_deconfigure
```

- d. As `root`, start ACFS kernel modules:

```
# $ORACLE_HOME/bin/acfsload start
```

- e. As `root`, start the Oracle Clusterware stack on the node:

```
# $ORACLE_HOME/bin/crsctl start crs
```

- f. As the Oracle Grid Infrastructure owner verify the status of Oracle ASMFD:

```
$ $ORACLE_HOME/bin/asmcmd afd_state
```

4. As the Oracle Grid Infrastructure owner update the Oracle ASM discovery string to remove the Oracle ASMFD path:

```
$ $ORACLE_HOME/bin/asmcmd dsget
```

```
$ $ORACLE_HOME/bin/asmcmd dsset old_diskstring
```

Check the current value of the Oracle ASM disk discovery string before updating the value. The *old_diskstring* value is the old disk discovery string value before updating with the AFD: (Oracle ASMFD) paths.

See Also:

- [Updating the Oracle ASM ASM_DISKSTRING Parameter for Oracle ASM Filter Driver Disks](#) (page 20-11) for information about updating the Oracle ASM disk discovery string
 - [Updating the Oracle ASM Filter Driver AFD_DISKSTRING Parameter](#) (page 20-10) for information about updating the Oracle ASM Filter Driver discovery disk discovery string
 - [acfsload](#) (page 17-6) for information about the `acfsload` command
 - [Determining Whether Oracle ASM Filter Driver Has Been Configured](#) (page 20-10) for information about checking on the state of the Oracle ASM Filter Driver
-

20.10.2 Deconfiguring Oracle ASM Filter Driver in an Oracle Grid Infrastructure Standalone (Oracle Restart) Environment

Perform the following steps to deconfigure Oracle ASM Filter Driver in an Oracle Grid Infrastructure standalone environment:

1. Update the Oracle ASM disk discovery string to enable Oracle ASM to discover disk devices after Oracle ASMFD is deconfigured.

2. Log in as the `root` user and stop Oracle Grid Infrastructure for a standalone server using the following command:

```
# $ORACLE_HOME/bin/crsctl stop has
```

If the previous command returns an error, then use the following command:

```
# $ORACLE_HOME/bin/crsctl stop has -f
```

3. As `root`, stop Oracle ACFS kernel modules to ensure the most reliable shutdown:

```
# $ORACLE_HOME/bin/acfsload stop
```

4. As `root`, deconfigure Oracle ASMFD:

```
# $ORACLE_HOME/bin/asmcmd afd_deconfigure
```

5. As `root`, start ACFS kernel modules:

```
# $ORACLE_HOME/bin/acfsload start
```

6. As `root`, start Oracle Grid Infrastructure for a standalone server:

```
# $ORACLE_HOME/bin/crsctl start has
```

7. As the Oracle Grid Infrastructure standalone server owner verify the Oracle ASMFD status:

```
$ $ORACLE_HOME/bin/asmcmd afd_state
```

8. As the Oracle Grid Infrastructure standalone server owner update the Oracle ASM disk discovery string to remove the Oracle ASMFD paths:

```
$ $ORACLE_HOME/bin/asmcmd dsget
```

```
$ $ORACLE_HOME/bin/asmcmd dsset old_diskstring
```

Check the current value of the Oracle ASM disk discovery string before updating the value. The *old_diskstring* value is the old disk discovery string value before updating with the AFD: (Oracle ASMFd) paths.

See Also:

- [Updating the Oracle ASM ASM_DISKSTRING Parameter for Oracle ASM Filter Driver Disks](#) (page 20-11) for information about updating the Oracle ASM disk discovery string
 - [Updating the Oracle ASM Filter Driver AFD_DISKSTRING Parameter](#) (page 20-10) for information about updating the Oracle ASM Filter Driver discovery disk discovery string
 - [acfsload](#) (page 17-6) for information about the `acfsload` command
 - [Determining Whether Oracle ASM Filter Driver Has Been Configured](#) (page 20-10) for information about checking on the state of the Oracle ASM Filter Driver
-
-

20.11 Oracle ASM Filter Driver Commands

This section describes the Oracle ASM Filter Driver (ASMFd) commands that are used during installation to manage Oracle ASM Filter Driver. These commands are located in the `/bin` directory of the Oracle Grid Infrastructure home.

- [afdload](#) (page 20-15)
- [afddriverstate](#) (page 20-16)

20.11.1 afdload

Purpose

`afdload` loads or unloads the Oracle ASM Filter Driver.

Syntax

```
afdload { start | stop } [ -s ]
```

[Table 20-1](#) (page 20-15) contains the options available with the `afdload` command.

Table 20-1 Options for the `afdload` command

Option	Description
<code>start</code>	Loads the Oracle ASM Filter Driver.
<code>stop</code>	Unloads the Oracle ASM Filter Driver.
<code>-s</code>	Operate in silent mode.

Description

You can use `afdload` to manually load or unload the Oracle ASM Filter Driver.

Before unloading driver with the `stop` option, you must dismount disk groups and shut down Oracle ASM.

`root` or administrator privilege is required to run `afdload`.

Examples

The following is an example of the use of `afdload` to stop (unload) the Oracle ASM Filter Driver.

```
# afdload stop
```

20.11.2 afddriverstate**Purpose**

`afddriverstate` provides information on the current state of the Oracle ASM Filter Driver.

Syntax

```
afddriverstate [-orahome ORACLE_HOME ]
               { installed | loaded | version | supported } [-s]
```

[Table 20-2](#) (page 20-16) contains the options available with the `afddriverstate` command.

Table 20-2 Options for the `afddriverstate` command

Option	Description
<code>-orahome ORACLE_HOME</code>	Specifies the Oracle Grid Infrastructure home in which the user has permission to execute the <code>afddriverstate</code> command.
<code>installed</code>	Determines whether Oracle ASM Filter Driver is installed on the system.
<code>loaded</code>	Determines whether the Oracle ASM Filter Driver is loaded in memory.
<code>version</code>	Reports the currently installed version of the Oracle ASM Filter Driver system software.
<code>supported</code>	Reports whether the system is a supported kernel for Oracle ASM Filter Driver.
<code>-s</code>	Specifies silent mode when running the command.

Description

You can use `afddriverstate` to display detailed information on the current state of the Oracle ASM Filter Driver.

Examples

The following is an example of the use of `afddriverstate`.

```
$ afddriverstate version
AFD-9325:    Driver operating system kernel version =
2.6.39-400.3.0.el6uek.x86_64(x86_64).
AFD-9326:    Driver Oracle version = 150728.
AFD-9212:    Driver build version = MAIN.
```

Creating Database Files in Oracle ASM Using SQL*Plus

Oracle Automatic Storage Management (Oracle ASM) can be deployed using SQL*Plus.

This section contains the following topics:

- [Creating a Database that Uses Oracle ASM](#) (page 21-1)
- [Creating Tablespaces in Oracle ASM](#) (page 21-2)
- [Creating Redo Logs in Oracle ASM](#) (page 21-2)
- [Creating Control Files in Oracle ASM](#) (page 21-3)
- [Creating Archive Log Files in Oracle ASM](#) (page 21-4)

See Also:

Oracle Database Administrator's Guide for information about using Oracle Managed Files (OMF)

21.1 Creating a Database that Uses Oracle ASM

The recommended method of creating your database is to use 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 example shows the use of the `CREATE DATABASE` statement, where database files are created and managed automatically by Oracle ASM. 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 can be run at the SQL prompt:

```
SQL> CREATE DATABASE orcl;
```

21.2 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](#) (page 5-19)". 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 due to 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`, then you can create a tablespace `tblspace` in that disk group with the following SQL statement:

```
SQL> 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:

```
SQL> CREATE TABLESPACE tblspace;
```

21.3 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 file names for log file members are automatically generated. All partially created redo log files, created because 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:

```
SQL> 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`. These 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 can be run at the SQL prompt:

```
SQL> ALTER DATABASE ADD LOGFILE;
```

21.4 Creating Control Files in Oracle ASM

Control files can be explicitly created in multiple disk groups. The file names for control files are automatically generated.

If you must specify a control file by name, then you can use alias file names 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 file names. 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](#) (page 21-3)
- [Creating Control Files in Oracle ASM: Example 2](#) (page 21-4)

21.4.1 Creating Control Files in Oracle ASM: Example 1

The `CREATE CONTROLFILE` SQL statement in [Example 21-1](#) (page 21-3) 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`:

Example 21-1 *Creating control files in Oracle ASM using NORESETLOGS*

```
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
;
```

21.4.2 Creating Control Files in Oracle ASM: Example 2

Example 21-2 (page 21-4) is a CREATE CONTROLFILE SQL statement for a database with data files, but uses a RESETLOGS clause, which uses the creation form for log files:

Example 21-2 Creating control files in Oracle ASM using RESETLOGS

```
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
;
```

21.5 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 file names are created in default subdirectories with unique names. However, all LOG_ARCHIVE_DEST_n parameters must have a unique string. To enter the same disk group more than once, subdirectories must be used. For example, you can set LOG_ARCHIVE_DEST_n as follows:

```
LOG_ARCHIVE_DEST_1 = 'LOCATION=+FRA/orcl/arc_1'
LOG_ARCHIVE_DEST_2 = 'LOCATION=+FRA/orcl/arc_2'
...
```

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 Oracle Database under default subdirectories. If LOG_ARCHIVE_DEST is set to a directory in a disk group, LOG_ARCHIVE_FORMAT has its normal semantics for the file alias. The underlying file is still generated in the default directories for the disk group with a system generated name.

You can also set the `DB_RECOVERY_FILE_DEST` initialization parameter to a disk group destination. In [Example 21-3](#) (page 21-5), archive log files are generated with `ARCHIVELOG` mode enabled, `DB_RECOVERY_FILE_DEST` set to the `+FRA` disk group, and `ORCL` is the value of the `DB_UNIQUE_NAME` parameter.

See Also:

Oracle Database Administrator's Guide for information about managing archived redo log files

Example 21-3 Archived log files in an Oracle ASM disk group

```
SQL> SHOW PARAMETER DB_RECOVERY_FILE_DEST
```

NAME	TYPE	VALUE
db_recovery_file_dest	string	+FRA

```
SQL> SELECT SUBSTR(name,1,64) name FROM V$ARCHIVED_LOG;
```

```
NAME
-----
+FRA/ORCL/ARCHIVELOG/2013_06_24/thread_1_seq_112.260.818960623
+FRA/ORCL/ARCHIVELOG/2013_06_24/thread_1_seq_113.261.818978423
+FRA/ORCL/ARCHIVELOG/2013_06_25/thread_1_seq_114.262.818993031
+FRA/ORCL/ARCHIVELOG/2013_06_25/thread_1_seq_115.263.819032439
+FRA/ORCL/ARCHIVELOG/2013_06_25/thread_1_seq_116.264.819047121
+FRA/ORCL/ARCHIVELOG/2013_06_25/thread_1_seq_117.265.819064815
+FRA/ORCL/ARCHIVELOG/2013_06_26/thread_1_seq_118.266.819079417
...
```

Glossary

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.

Oracle ASM Filter Driver is the recommended replacement for ASMLIB. For information about Oracle ASM Filter Driver, refer to "[Oracle ASM Filter Driver](#) (page 20-1)".

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.

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 optical fiber 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 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.

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 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.

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 are 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**.

Symbols

.ACFS directory
Oracle ACFS snapshots, [11-11](#), [16-127](#)

A

accelerator volume

Oracle ADVM, [11-43](#), [16-28](#)

access control

disk groups, [4-61](#)

ACCESS_CONTROL.ENABLED disk group attribute,
[4-4](#), [4-63](#)

ACCESS_CONTROL.UMASK disk group attribute,
[4-4](#), [4-63](#)

ACFS

See Oracle Automatic Storage Management Cluster
File System

acfschkdisk command, [16-50](#)

acfsdismount command, [16-52](#)

acfsdriverstate

Oracle ACFS driver command, [17-7](#)

acfsfileid_lookup

Oracle ACFS Plug-in APIs, [17-9](#)

acfsformat command, [16-53](#)

acfsgettag

Oracle ACFS tagging API, [17-14](#)

acfslisttags

Oracle ACFS tagging API, [17-15](#)

acfsload

Oracle ACFS driver command, [17-6](#)

Oracle ASM Filter Driver command, [20-15](#)

AcfsMaxCachedFiles

Oracle ACFS tunable parameter, [16-168](#)

AcfsMaxOpenFiles

Oracle ACFS tunable parameter, [16-168](#)

acfsmountvol command, [16-54](#)

acfsplugin_metrics

Oracle ACFS Plug-in APIs, [17-9](#)

acfsremovetag

Oracle ACFS tagging API, [17-16](#)

acfssettag

Oracle ACFS tagging API, [17-17](#)

acfsutil accel replace command, [16-143](#)

acfsutil acfsdbg command, [17-19](#)

acfsutil audit archive command, [16-81](#)

acfsutil audit disable command, [16-82](#)

acfsutil audit enable command, [16-82](#)

acfsutil audit info command, [16-83](#)

acfsutil audit init command, [16-84](#)

acfsutil audit purge command, [16-85](#)

acfsutil audit read command, [16-85](#)

acfsutil blog command, [17-22](#)

acfsutil command-line tools

running on Windows, [16-2](#)

acfsutil compat get command, [16-144](#)

acfsutil compat set command, [16-145](#)

acfsutil compress copy command, [16-140](#)

acfsutil compress info command, [16-141](#)

acfsutil compress off command, [16-142](#)

acfsutil compress on command, [16-142](#)

acfsutil defrag dir command, [16-146](#)

acfsutil defrag file command, [16-146](#)

acfsutil detach command, [16-56](#)

acfsutil dumpstate command, [17-23](#)

acfsutil encr command, [16-121](#)

acfsutil encr info command, [16-118](#)

acfsutil encr init command, [16-119](#)

acfsutil encr off command, [16-120](#)

acfsutil encr rekey command, [16-122](#)

acfsutil encr set command, [16-123](#)

acfsutil info file command, [16-147](#)

acfsutil info fs command, [16-149](#)

acfsutil info id command, [16-154](#)

acfsutil keystore migrate command, [16-125](#)

acfsutil log command, [17-24](#)

acfsutil meta command, [17-26](#)

acfsutil plogconfig command, [17-28](#)

acfsutil plugin disable command, [16-155](#)

acfsutil plugin enable command, [16-156](#)

acfsutil plugin info command, [16-159](#)

acfsutil registry command, [16-160](#)

acfsutil repl bg command, [16-64](#)

acfsutil repl info command, [16-67](#)

acfsutil repl info compare, [16-65](#)

acfsutil repl init command, [16-69](#)

- acfsutil repl pause command, [16-72](#)
- acfsutil repl resume command, [16-73](#)
- acfsutil repl sync command, [16-73](#)
- acfsutil repl terminate command, [16-75](#)
- acfsutil repl trace command, [16-76](#)
- acfsutil repl update command, [16-76](#)
- acfsutil repl upgrade command, [16-78](#)
- acfsutil rmfs command, [16-162](#)
- acfsutil scrub command, [16-163](#)
- acfsutil sec admin add command, [16-87](#)
- acfsutil sec admin password command, [16-88](#)
- acfsutil sec admin remove command, [16-88](#)
- acfsutil sec batch command, [16-89](#)
- acfsutil sec disable command, [16-90](#)
- acfsutil sec enable command, [16-91](#)
- acfsutil sec info command, [16-91](#)
- acfsutil sec info file command, [16-92](#)
- acfsutil sec init command, [16-93](#)
- acfsutil sec load command, [16-95](#)
- acfsutil sec prepare command, [16-96](#)
- acfsutil sec realm add command, [16-98](#)
- acfsutil sec realm audit disable command, [16-102](#)
- acfsutil sec realm audit enable command, [16-103](#)
- acfsutil sec realm audit info command, [16-105](#)
- acfsutil sec realm clone command, [16-105](#)
- acfsutil sec realm create command, [16-107](#)
- acfsutil sec realm delete command, [16-108](#)
- acfsutil sec realm destroy command, [16-109](#)
- acfsutil sec rule clone command, [16-109](#)
- acfsutil sec rule create command, [16-110](#)
- acfsutil sec rule destroy command, [16-112](#)
- acfsutil sec rule edit command, [16-112](#)
- acfsutil sec ruleset clone command, [16-113](#)
- acfsutil sec ruleset create command, [16-114](#)
- acfsutil sec ruleset destroy command, [16-115](#)
- acfsutil sec ruleset edit command, [16-116](#)
- acfsutil sec save command, [16-116](#)
- acfsutil size command
 - automatic and manual resizing, [16-164](#)
- acfsutil snap convert command, [16-126](#)
- acfsutil snap create command, [16-127](#)
- acfsutil snap delete command, [16-130](#)
- acfsutil snap duplicate apply command, [16-131](#)
- acfsutil snap duplicate create command, [16-132](#)
- acfsutil snap info command, [16-134](#)
- acfsutil snap quota command, [16-136](#)
- acfsutil snap remaster command, [16-138](#)
- acfsutil snap rename command, [16-139](#)
- acfsutil tag info command, [16-60](#)
- acfsutil tag set command, [16-61](#)
- acfsutil tag unset command, [16-63](#)
- acfsutil tune command, [16-167](#)
- active session history sampling
 - Oracle Automatic Storage Management, [3-24](#)
- adding aliases
 - Oracle ASM, [5-14](#)
- adding disks
 - disk groups, [4-22](#)
- adding disks to a disk group, [1-12](#)
- adding disks to a disk group with Oracle Enterprise Manager, [7-7](#)
- adding disks to an Oracle ASM disk group
 - Oracle ASM Configuration Assistant, [9-10](#)
- adding objects to security realms
 - Oracle ACFS, [16-9](#)
- adding security administrators
 - Oracle ACFS file systems, [16-9](#)
- adding templates to a disk group, [5-22](#)
- administering
 - disk groups, [4-1](#)
 - Oracle ACFS, [11-15](#)
- advanced topics
 - Oracle ACFS, [17-1](#)
- advmutil canonical command, [16-168](#)
- advmutil dismount command, [16-56](#)
- advmutil list command, [16-58](#)
- advmutil mount command, [16-58](#)
- advmutil tune command, [16-169](#)
- advmutil volinfo command, [16-170](#)
- afd_configure command
 - ASMCMD utility, [10-33](#)
- afd_deconfigure command
 - ASMCMD utility, [10-34](#)
- afd_di command
 - ASMCMD utility, [10-35](#)
- afd_dsget command
 - ASMCMD utility, [10-36](#)
- afd_dsset command
 - ASMCMD utility, [10-36](#)
- afd_filter command
 - ASMCMD utility, [10-36](#)
- afd_label command
 - ASMCMD utility, [10-37](#)
- afd_lsdsdsk command
 - ASMCMD utility, [10-38](#)
- afd_lsbl command
 - ASMCMD utility, [10-38](#)
- afd_refresh command
 - ASMCMD utility, [10-39](#)
- afd_scan command
 - ASMCMD utility, [10-40](#)
- afd_state command
 - ASMCMD utility, [10-40](#)
- afd_unlabel command
 - ASMCMD utility, [10-41](#)
- afddriverstate
 - Oracle ASM Filter Driver command, [20-16](#)
- AIX
 - Oracle ACFS command-line tools, [16-41](#)
- alias file names
 - Oracle ASM, [5-6](#)
- aliases
 - adding for an Oracle ASM file name, [5-14](#)
 - dropping from a disk group, [5-15](#)

ALL command rule, [16-98](#)

allocation units

- different sizes enabled by compatibility attributes, [4-58](#)
- Oracle ASM, [1-7](#)
- size for disk groups, [4-4](#)

ALTER DISKGROUP command, [4-19](#)

altering disk groups

- replacing disks, [4-24](#)

altering sector size

- enabled by compatibility attributes, [4-58](#)

amdu_extract command

- ASMCMD utility, [10-42](#)

APPENDFILE command rule, [16-98](#)

appliance mode

- enabled by compatibility attributes, [4-58](#)

archive log files

- creating in Oracle ASM, [21-4](#)

ASM

- See Oracle Automatic Storage Management

ASM SPFILE in a disk group

- enabled by compatibility attributes, [4-58](#)

ASM_DISKGROUPS initialization parameter

- Oracle ASM instance, [3-9](#)

ASM_DISKSTRING initialization parameter

- discovering disks, [1-11](#)
- Oracle ASM instance, [3-10](#)

ASM_IO_PROCESSES initialization parameter

- Oracle IO Server, [3-11](#)

ASM_POWER_LIMIT initialization parameter

- Oracle ASM instance, [3-11](#)
- rebalancing process, [3-11](#)
- values enabled by compatibility attributes, [4-58](#)

ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter

- Oracle ASM instance, [3-12](#)

asmadmin group, [3-29](#)

ASMCA

- See Oracle ASM Configuration Assistant

ASMCMD utility

- about, [10-1](#)
- afd_configure command, [10-33](#)
- afd_deconfigure command, [10-34](#)
- afd_di command, [10-35](#)
- afd_dsget command, [10-36](#)
- afd_dsset command, [10-36](#)
- afd_filter command, [10-36](#)
- afd_label command, [10-37](#)
- afd_ksdsk command, [10-38](#)
- afd_kslbl command, [10-38](#)
- afd_refresh command, [10-39](#)
- afd_scan command, [10-40](#)
- afd_state command, [10-40](#)
- afd_unlabel command, [10-41](#)
- aliases, and, [10-4](#)
- amdu_extract command, [10-42](#)
- cd command, [10-43](#)
- ASMCMD utility (*continued*)
 - chdg command, [10-54](#)
 - chfg command, [10-79](#)
 - chgrp command, [10-91](#)
 - chkdg command, [10-57](#)
 - chmod command, [10-92](#)
 - chown command, [10-93](#)
 - chtmpl command, [10-87](#)
 - client cluster management commands, [10-30](#)
 - command-line utility, [10-1](#)
 - commands for Oracle ACFS and Oracle ADVM, [15-1](#)
 - cp command, [10-44](#)
 - directories, and, [10-3](#)
 - disk group management commands, [10-53](#)
 - displaying the version number, [10-9](#)
 - dropdg command, [10-57](#)
 - dsget command, [10-13](#)
 - dsset command, [10-13](#)
 - du command, [10-45](#)
 - exit command, [10-9](#)
 - file access control commands, [10-90](#)
 - file group commands, [10-78](#)
 - file management commands, [10-42](#)
 - find command, [10-46](#)
 - groups command, [10-94](#)
 - grpmod command, [10-95](#)
 - help command, [10-11](#)
 - including directory in prompt, [10-9](#)
 - instance management commands, [10-11](#)
 - iostat command, [10-58](#)
 - ls command, [10-47](#)
 - lsattr command, [10-60](#)
 - lscc command, [10-30](#)
 - lsct command, [10-14](#)
 - lsdg command, [10-61](#)
 - ksdsk command, [10-63](#)
 - ksgrp command, [10-96](#)
 - ksod command, [10-65](#)
 - ksop command, [10-49](#)
 - ksop command, [10-15](#)
 - kspwusr command, [10-15](#)
 - ksmpl command, [10-88](#)
 - ksusr command, [10-96](#)
 - md_backup command, [10-66](#)
 - md_restore command, [10-67](#)
 - mkalias command, [10-50](#)
 - mkcc command, [10-31](#)
 - mkdg command, [10-69](#)
 - mkdir command, [10-51](#)
 - mkfg command, [10-83](#)
 - mkgrp command, [10-97](#)
 - ksmpl command, [10-89](#)
 - ksusr command, [10-98](#)
 - mount command, [10-71](#)
 - offline command, [10-72](#)

ASMCMD utility (*continued*)

- online command, [10-73](#)
- Oracle ADVM volume management commands, [15-1](#)
- Oracle ASM Filter Driver management
 - commands, [10-32](#)
- orapwusr command, [10-16](#)
- passwd command, [10-98](#)
- preparing to run, [10-6](#)
- pwcopy command, [10-17](#)
- pwcreate command, [10-18](#)
- pwd command, [10-51](#)
- pwdelete command, [10-19](#)
- pwget command, [10-19](#)
- pwmove command, [10-20](#)
- pwset command, [10-21](#)
- rebal command, [10-74](#)
- remap command, [10-75](#)
- rm command, [10-52](#)
- rmalias command, [10-53](#)
- rmcc command, [10-32](#)
- rmgrp command, [10-99](#)
- rmtmpl command, [10-90](#)
- rmusr command, [10-99](#)
- rpusr command, [10-100](#)
- running in a database home, [10-6](#)
- running in an Oracle grid infrastructure home, [10-6](#)
- running in interactive mode, [10-7](#)
- running in noninteractive mode, [10-10](#)
- running without any options, [10-8](#)
- setattr command, [10-75](#)
- showclustermode command, [10-22](#)
- showclusterstate command, [10-22](#)
- showpatches command, [10-22](#)
- showversion command, [10-23](#)
- shutdown command, [10-23](#)
- spbackup command, [10-24](#)
- spcopy command, [10-25](#)
- specifying the verbose mode, [10-9](#)
- spget command, [10-27](#)
- spmmove command, [10-27](#)
- spset command, [10-28](#)
- stamp command, [10-76](#)
- stamplist command, [10-77](#)
- startup command, [10-29](#)
- syntax for starting in interactive mode, [10-8](#)
- SYSASM privilege, [10-9](#)
- system-generated file names, and, [10-3](#)
- template management commands, [10-87](#)
- type of connection, [10-9](#)
- types of commands, [10-2](#)
- umount command, [10-78](#)
- using quotes for wildcards in noninteractive mode, [10-10](#)
- using with absolute and relative paths, [10-4](#)

ASMCMD utility (*continued*)

- using with wildcard characters, [10-5](#)
- volcreate command, [15-2](#)
- voldelete command, [15-4](#)
- voldisable command, [15-5](#)
- volenable command, [15-6](#)
- volinfo command, [15-6](#)
- volresize command, [15-7](#)
- volset command, [15-8](#)
- volstat command, [15-9](#)

asmdba group, [3-29](#)

ASMF

- configuring, [20-3](#)
- configuring after installation, [20-5](#)
- configuring during installation, [20-4](#)

asmoper group, [3-29](#)

ASMSNMP user

- for monitoring the Oracle ASM instance, [3-28](#)
- Oracle ASM, [9-2](#)
- Oracle ASM in Oracle Enterprise Manager, [7-3](#)

attributes

- disk group compatibility, [4-51](#)
- disk groups, [4-2](#)
- DISK_REPAIR_TIME, [4-43](#)
- for disk groups, [4-12](#)
- in templates for Oracle Automatic Storage Management, [5-20](#)
- list of for disk groups, [4-3](#)
- setting for disk groups, [4-2](#)
- viewing for disk groups, [4-2](#)

AU_SIZE disk group attribute

- different sizes enabled by compatibility attributes, [4-58](#)

audit trail

- Oracle ASM, [3-36](#)

auditing

- Oracle ACFS audit trail file, [11-31](#)
- Oracle ACFS command-line tools, [16-81](#)
- Oracle ACFS evaluation result events, [11-35](#)
- Oracle ACFS file access events, [11-33](#)
- Oracle ACFS privilege use events, [11-34](#)
- Oracle ACFS security and encryption, [11-29](#)
- Oracle ASM Cluster File System, [11-30](#)

authentication

- operating system with Oracle ASM, [3-32](#)
- Oracle ASM, [3-27](#)
- password file with Oracle ASM, [3-32](#)

Automatic Storage Management

- home page on Oracle Enterprise Manager, [13-1](#)

B

- backing up Oracle ASM files
 - Oracle Enterprise Manager, [7-22](#)
- backup and restore
 - Oracle ACFS, [11-14](#)
- bad block recovery in Oracle ASM disks

bad block recovery in Oracle ASM disks (*continued*)
Oracle Enterprise Manager, 7-23

C

capacity

managing in disk groups, 4-35

cd command

ASMCMD utility, 10-43

CELL.SMART_SCAN_CAPABLE disk group attribute,
4-4

CELL.SPARSE_DG disk group attribute, 4-4

changing tag names

Oracle ACFS, 16-16

chdg command

ASMCMD utility, 10-54

checking Oracle ASM disk group integrity

Oracle Enterprise Manager, 7-22

chfg command

ASMCMD utility, 10-79

chgrp command

ASMCMD utility, 10-91

CHGRP command rule, 16-98

chkdsk command

ASMCMD utility, 10-57

chmod command

ASMCMD utility, 10-92

CHMOD command rule, 16-98

chown command

ASMCMD utility, 10-93

CHOWN command rule, 16-98

chtmpl command

ASMCMD utility, 10-87

client cluster management commands

ASMCMD utility, 10-30

Oracle ASM remote instances, 10-30

command rules

ALL, 16-98

APPENDFILE, 16-98

CHGRP, 16-98

CHMOD, 16-98

CHOWN, 16-98

CREATEFILE, 16-98

DELETEFILE, 16-98

EXTEND, 16-98

IMMUTABLE, 16-98

LINKFILE, 16-98

MKDIR, 16-98

MMAPREAD, 16-98

MMAPWRITE, 16-98

OPENFILE, 16-98

Oracle ACFS security, 16-98

OVERWRITE, 16-98

READ, 16-98

REaddir, 16-98

RENAME, 16-98

RMDIR, 16-98

command rules (*continued*)

SYMLINK, 16-98

TRUNCATE, 16-98

WRITE, 16-98

command-line tools

Oracle ACFS, 16-1

Oracle ACFS auditing, 16-81

Oracle ACFS compression, 16-140

Oracle ACFS encryption, 16-117

Oracle ACFS for auditing, 16-81

Oracle ACFS security, 16-86

Oracle ACFS snapshots, 16-125

command-line tools for replication

Oracle ACFS, 16-63

command-line tools for tagging

Oracle ACFS, 16-60

command-line utilities

Oracle ACFS, 16-143

commands

Oracle ASM Filter Driver command, 20-15

compatibility

disk groups, 4-50, 7-5

Oracle ASM and database instances, 3-1

reverting disk groups to previous settings, 4-55

COMPATIBLE

database initialization parameter, 4-52

COMPATIBLE.ADVm disk group attribute

default value, 4-52

COMPATIBLE.ASM disk group attribute

default value, 4-52

COMPATIBLE.RDBMS disk group attribute

default value, 4-52

components

Oracle ASM, 1-3

concepts

Oracle ASM, 1-1

configuration page

Oracle ASM on Oracle Enterprise Manager, 7-2

configuring an Oracle ASM instance

Oracle ASM Configuration Assistant, 9-6

Oracle Enterprise Manager, 7-2

configuring parameters for an Oracle ASM instance

Oracle ASM Configuration Assistant, 9-7

configuring snapshot-based replication, 18-1

configuring ssh with snapshot-based replication, 18-1

content type of a disk group

enabled by compatibility attributes, 4-58

CONTENT.CHECK disk group attribute, 4-5

CONTENT.TYPE disk group attribute, 4-5

control file

creating in Oracle ASM, 21-3

converting snapshots

enabled by compatibility attributes, 4-59

cp command

ASMCMD utility, 10-44

CREATEFILE command rule, 16-98

createFileGroups, 9-12

- createLabel, [9-12](#)
- createQuotaGroups, [9-13](#)
- creating a database home
 - with Oracle ASMCA, [14-5](#)
- creating a database using Oracle ASM, [21-1](#)
- creating a new directory
 - Oracle ASM, [5-16](#)
- creating an Oracle ACFS file system
 - Oracle ASM Configuration Assistant, [14-8](#)
- creating an Oracle ACFS file system snapshot
 - Oracle ASM Configuration Assistant, [14-7](#)
- creating an Oracle ADVM volume
 - Oracle ASM Configuration Assistant, [14-8](#)
- creating an Oracle ASM disk group
 - Oracle ASM Configuration Assistant, [9-9](#)
- creating database files in Oracle ASM, [21-1](#)
- creating disk groups, [4-9](#), [4-10](#)
- creating disk groups and adding disks, [4-10](#)
- creating disk groups and attributes, [4-10](#)
- creating disk groups and failure group types, [4-10](#)
- creating disk groups and failure groups, [4-10](#)
- creating disk groups and redundancy level, [4-10](#)
- creating disk groups in a new installation, [4-14](#)
- creating files in the database
 - Oracle ASM, [5-9](#)
- creating files using a default file location
 - Oracle ASM, [5-9](#)
- creating from existing snapshots
 - enabled by compatibility attributes, [4-59](#)
- creating Oracle ACFS file systems, [16-5](#)
- creating security realms
 - Oracle ACFS, [16-9](#)
- creating security rule sets
 - Oracle ACFS, [16-9](#)
- creating security rules
 - Oracle ACFS, [16-9](#)

D

- data files
 - moving between disk groups, [5-11](#)
- data files with alias names
 - Oracle ASM, [5-7](#)
- database files
 - Oracle ACFS on Oracle Exadata, [11-6](#)
- DB_CACHE_SIZE initialization parameter
 - Oracle ASM instance, [3-13](#)
- default configuration
 - Oracle ADVM, [15-2](#), [17-4](#)
- DELETEFILE command rule, [16-98](#)
- deleting an Oracle ACFS file system snapshot
 - Oracle ASM Configuration Assistant, [14-7](#)
- deleting Oracle ASM instances
 - Oracle ASM Configuration Assistant, [9-8](#)
- deregistering file systems
 - Oracle ACFS, [16-23](#)
- determining the volume device name

- determining the volume device name (*continued*)
 - Oracle ACFS, [16-5](#)
- diagnostic commands
 - Oracle ACFS, [17-18](#)
- diagnostic directory
 - Oracle ASM instance, [3-14](#)
- DIAGNOSTIC_DEST initialization parameter
 - Oracle ASM instance, [3-14](#)
- directories
 - ACFS for snapshots, [11-11](#), [16-127](#)
 - creating disk group, [5-16](#)
 - dropping disk group, [5-17](#)
 - managing disk group, [5-16](#)
 - Oracle ASM, [5-16](#)
 - renaming disk group, [5-16](#)
- disabling a volume
 - Oracle ACFS, [16-24](#)
- disk discovery
 - about, [4-33](#)
 - improving time, [4-35](#)
 - in Oracle ASM, [4-33](#)
 - Oracle ASM, [1-11](#)
 - rules, [4-34](#)
- disk failures
 - in Oracle Automatic Storage Management, [4-41](#)
- disk group compatibility
 - configuring with Oracle Enterprise Manager, [7-13](#)
 - considerations about file sizes, [4-60](#)
 - considerations in a replicated environments, [4-60](#)
 - considerations when setting, [4-60](#)
 - Oracle ACFS features enabled by settings, [4-57](#)
 - Oracle ASM features enabled by settings, [4-57](#)
 - setting, [4-53](#)
 - setting in Oracle Enterprise Manager, [7-5](#)
 - setting with ALTER DISKGROUP, [4-54](#)
 - setting with CREATE DISKGROUP, [4-54](#)
 - valid combinations, [4-53](#)
 - viewing compatibility settings, [4-55](#)
- disk group management commands
 - ASMCMD utility, [10-53](#)
- disk group properties
 - configuring with Oracle Enterprise Manager, [7-13](#)
- disk group repair time
 - configuring with Oracle Enterprise Manager, [7-14](#)
- disk group smart scan compatibility
 - configuring with Oracle Enterprise Manager, [7-14](#)
- disk groups
 - access control, [4-61](#)
 - ACCESS_CONTROL.ENABLED attribute, [4-63](#)
 - ACCESS_CONTROL.UMASK attribute, [4-63](#)
 - adding disks, [4-22](#)
 - adding disks to with Oracle ASM Configuration Assistant, [9-10](#)
 - adding disks to with Oracle Enterprise Manager, [7-7](#)
 - adding templates to, [5-22](#)
 - administering, [4-1](#)

- disk groups (*continued*)
 - allocation unit size, [4-4](#)
 - altering membership of, [4-19](#)
 - attributes, [4-2](#), [4-12](#)
 - checking internal consistency, [4-68](#)
 - compatibility, [4-50](#)
 - compatibility attributes, [4-51](#)
 - converting to flex disk groups, [4-73](#)
 - creating, [4-9](#)
 - creating and adding disks, [4-10](#)
 - creating and attributes, [4-10](#)
 - creating and failure group types, [4-10](#)
 - creating and failure groups, [4-10](#)
 - creating and redundancy level, [4-10](#)
 - creating for a new installation, [4-14](#)
 - creating Oracle Enterprise Manager, [7-5](#)
 - creating with Oracle ASM Configuration Assistant, [9-9](#)
 - creating with SQL*Plus, [4-10](#)
 - determining the number of, [4-48](#)
 - directories, [5-16](#)
 - discovering disks, [4-33](#)
 - dismounting, [1-12](#)
 - dismounting with Oracle Enterprise Manager, [7-12](#)
 - dropping, [4-69](#)
 - dropping disks from, [4-25](#)
 - dropping disks with Oracle Enterprise Manager, [7-9](#)
 - dropping in Oracle Enterprise Manager, [7-11](#)
 - failure groups, [4-38](#), [4-40](#)
 - flex disk groups, [4-73](#)
 - list of attributes, [4-3](#)
 - logical sector size, [4-15](#)
 - managing a shared password file, [3-34](#)
 - managing capacity in, [4-35](#)
 - manually rebalancing, [4-28](#)
 - migrating to flex disk groups, [4-73](#)
 - mirroring, [4-38](#)
 - monitoring disk space usage in Oracle Enterprise Manager, [7-12](#)
 - mounting, [1-12](#)
 - mounting and dismounting, [4-67](#)
 - mounting at startup, [3-21](#)
 - mounting with force, [4-67](#)
 - mounting with Oracle Enterprise Manager, [7-12](#)
 - moving data files between, [5-11](#)
 - names, [4-10](#)
 - Oracle ASM, [1-6](#)
 - Oracle Cluster Registry (OCR), and, [4-18](#)
 - performance and scalability, [4-48](#)
 - performance characteristics when grouping disks, [4-48](#)
 - preferred read, [3-12](#)
 - quorum failure groups, [4-18](#)
 - rebalancing, [1-13](#)

- disk groups (*continued*)
 - redundancy, [4-38](#)
 - renamedg, [4-70](#)
 - renaming, [4-70](#)
 - renaming disks, [4-24](#)
 - replacing disks, [4-24](#)
 - resizing disks in, [4-27](#)
 - reverting compatibility, [4-55](#)
 - scrubbing, [4-32](#)
 - sector size, [4-15](#)
 - setting attributes, [4-2](#)
 - storage limits, [4-49](#)
 - tracking the replication status, [4-7](#)
 - undropping disks in, [4-28](#)
 - validations for logical sector size, [4-15](#)
 - validations for sector size, [4-15](#)
 - viewing attributes, [4-2](#)
 - voting files, and, [4-18](#)
 - when to use external, normal, or high redundancy, [4-42](#)
- disk partitions
 - storage resources, [2-1](#)
- disk regions
 - intelligent data placement, [4-26](#)
 - Oracle Automatic Storage Management, [4-26](#)
 - rebalance operations, and, [4-26](#)
- DISK_REPAIR_TIME disk group attribute
 - fast disk resync, [4-43](#)
- disks
 - adding to a disk group, [1-12](#)
 - discovering, [1-11](#), [4-33](#)
 - dropping from a disk group, [1-12](#)
 - Oracle ASM, [1-7](#)
 - regions, [4-26](#)
- dismounting disk groups, [1-12](#)
- dismounting file systems
 - Oracle ACFS, [16-24](#)
- dismounting or shutting down
 - Oracle ACFS, [11-22](#)
- displaying encryption information
 - Oracle ACFS, [16-14](#)
- displaying security information
 - Oracle ACFS, [16-9](#)
- displaying tagging information
 - Oracle ACFS, [16-16](#)
- driver commands
 - acfsdriverstate, [17-7](#)
 - acfsload, [17-6](#)
 - afddriverstate, [20-16](#)
 - afdload, [20-15](#)
 - Oracle ACFS, [17-6](#)
 - Oracle ASM Filter Driver, [20-15](#)
- driver model
 - Oracle ACFS, [11-10](#)
- drivers resource
 - Oracle ACFS, [17-4](#)
- dropdg command

- dropdg command (*continued*)
 - ASMCMD utility, [10-57](#)
- dropFileGroups, [9-13](#)
- dropLabel, [9-12](#)
- dropping a directory
 - Oracle ASM, [5-17](#)
- dropping aliases
 - Oracle ASM, [5-15](#)
- dropping aliases from a disk group, [5-15](#)
- dropping disk groups, [4-69](#)
- dropping disks from a disk group, [1-12](#), [4-25](#)
- dropping disks from disk groups in Oracle Enterprise Manager, [7-9](#)
- dropping files from a disk group
 - Oracle ASM, [5-15](#)
- dropping Oracle ASM templates, [5-23](#)
- dropQuotaGroups, [9-14](#)
- dsget command
 - ASMCMD utility, [10-13](#)
- dsset command
 - ASMCMD utility, [10-13](#)
- du command
 - ASMCMD utility, [10-45](#)
- dynamic views
 - Oracle ASM, [6-1](#)

E

- editDiskGroupAttributes, [9-11](#)
- enabling encryption
 - Oracle ACFS, [16-14](#)
- encrypting file systems
 - Oracle ACFS, [16-14](#)
- encrypting secure file systems
 - Oracle ACFS, [16-9](#)
- encryption
 - configuring with Oracle ASMCA, [14-4](#)
 - enabled by compatibility attributes, [4-59](#)
 - Oracle ACFS, [11-25](#)
 - Oracle ACFS and snapshots, [11-11](#)
 - Oracle ACFS command-line tools, [16-117](#)
- encryption keys
 - Oracle ACFS, [11-25](#)
- Enterprise Manager
 - See* Oracle Enterprise Manager
- Enterprise Manager Cloud Control
 - See* Oracle Enterprise Manager
- error handling
 - Oracle ACFS, [17-3](#)
- Exadata storage
 - enabled by compatibility attributes, [4-58](#)
- exit command
 - ASMCMD utility, [10-9](#)
- EXTEND command rule, [16-98](#)
- Extended Attributes
 - requirements for tagging, [11-39](#)
- extents
 - Oracle ASM, [1-7](#), [1-8](#)

- extents (*continued*)
 - variable size, [1-7](#), [1-8](#), [4-35](#), [4-58](#)
- external tables
 - and Oracle ACFS, [11-15](#)

F

- FAILGROUP_REPAIR_TIME disk group attribute
 - fast disk resync, [4-43](#)
- failure groups
 - about, [4-40](#)
 - disk failures, [4-41](#)
 - guidelines, [4-41](#)
 - how many to create, [4-42](#)
 - minimum recommended, [4-41](#)
 - mirroring, [4-38](#)
 - mirroring, and, [4-38](#)
 - Oracle ASM, [1-6](#)
 - preferred read, [4-46](#)
 - quorum type, [4-18](#)
 - simultaneous failure, [4-42](#)
 - when to use external, normal, or high redundancy, [4-42](#)
- fast mirror resync
 - enabled by compatibility attributes, [4-58](#)
 - managing with Oracle Enterprise Manager, [7-8](#), [7-9](#)
- file access and security
 - Oracle ACFS, [11-16](#)
- file access control
 - enabled by compatibility attributes, [4-58](#)
 - enabling with Oracle Enterprise Manager, [7-14](#)
 - managing with SQL statements, [4-64](#)
 - Oracle Automatic Storage Management, [4-61](#)
 - restricting access to files, [4-62](#)
 - setting up for files in a disk group, [4-62](#)
 - setting with SQL statements, [4-63](#)
- file management commands
 - ASMCMD utility, [10-42](#)
- file names
 - Oracle ASM, [5-3](#)
 - Oracle Automatic Storage Management, [5-3](#)
- file names in SQL statements
 - Oracle ASM, [5-10](#)
- file names with template form
 - Oracle ASM, [5-7](#)
- file templates
 - understanding, [1-11](#)
- file types supported
 - Oracle ASM, [5-1](#)
- files
 - Oracle ASM, [1-8](#)
 - templates, [1-11](#)
- filter driver
 - Oracle ASM, [20-2](#)
- find command
 - ASMCMD utility, [10-46](#)

- flex disk groups
 - Oracle ASM, [4-73](#)
- fsck command
 - AIX, [16-42](#)
 - Linux environments, [16-26](#)
 - Solaris, [16-34](#)
- fully qualified file names
 - Oracle ASM, [5-5](#)

G

- greater than 2 TB Oracle ASM disks
 - enabled by compatibility attributes, [4-58](#)
- grid infrastructure configuration
 - Oracle ACFS, [11-17](#)
- grid infrastructure installation
 - Oracle ACFS, [11-17](#)
- grid plug and play (GPnP)
 - retrieving the location of an ASM SPFILE, [10-27](#)
 - setting the location of the ASM SPFILE, [10-28](#)
- groups
 - OSASM, [3-27](#), [3-31](#), [3-32](#)
- groups command
 - ASMCMD utility, [10-94](#)
- grpmod command
 - ASMCMD utility, [10-95](#)

H

- help command
 - ASMCMD utility, [10-11](#)
- High Availability NFS for Oracle Grid Home
 - notes, [11-44](#)
 - Oracle ACFS, [11-44](#)
 - setting up for an Oracle ACFS file system, [11-44](#)

I

- I/O failure console messages
 - Oracle ACFS, [17-29](#)
- IDP.BOUNDARY disk group attribute, [4-6](#)
- IDP.TYPE disk group attribute, [4-6](#)
- IMMUTABLE command rule, [16-98](#)
- incomplete file name form
 - Oracle ASM, [5-8](#)
- incomplete file name template form
 - Oracle ASM, [5-8](#)
- individual file system resource
 - Oracle ACFS, [17-5](#)
- initialization parameter files
 - backing up, copying, and moving, [3-4](#)
 - creating an SPFILE in a disk group, [3-5](#)
 - Oracle ASM, [3-3](#)
 - search order for Oracle ASM, [3-3](#)
- initialization parameters
 - COMPATIBLE, [4-52](#)
 - database for use with Oracle ASM instance, [3-16](#)

- initialization parameters (*continued*)
 - Oracle ASM instance, [3-8](#)
 - Oracle Automatic Storage Management, [3-2](#)
 - setting for Oracle ASM instance, [3-7](#)
- initializing encryption on file systems
 - Oracle ACFS, [16-14](#)
- initializing security
 - Oracle ACFS file systems, [16-9](#)
- installing ssh and Cygwin for use with snapshot-based replication, [18-5](#)
- instance management commands
 - ASMCMD utility, [10-11](#)
- INSTANCE_TYPE initialization parameter
 - Oracle ASM instance, [3-14](#)
- instances
 - Oracle ASM, [1-3](#)
- intelligent data placement
 - disk regions, [4-26](#)
 - enabled by compatibility attributes, [4-58](#)
 - rebalance operations, and, [4-26](#)
- iostat command
 - ASMCMD utility, [10-58](#)

K

- keys
 - encryption, [11-25](#)

L

- LARGE_POOL_SIZE initialization parameter
 - Oracle ASM instance, [3-14](#)
- limits
 - Oracle ACFS, [17-1](#)
 - Oracle ADVM, [17-4](#)
- LINKFILE command rule, [16-98](#)
- Linux environments
 - Oracle ACFS command-line tools, [16-25](#)
- logical sector size
 - disk groups, [4-15](#)
 - validations for disk groups, [4-15](#)
- logical sector size of the Oracle ADVM volume
 - enabled by compatibility attributes, [4-59](#)
- logical unit number
 - storage resources, [2-1](#)
- logical volumes
 - storage resources, [2-1](#)
- LOGICAL_SECTOR_SIZE
 - enabled by compatibility attributes, [4-58](#)
- LOGICAL_SECTOR_SIZE disk group attribute, [4-7](#)
- ls command
 - ASMCMD utility, [10-47](#)
- lsattr command
 - ASMCMD utility, [10-60](#)
- lsc command
 - ASMCMD utility, [10-30](#)
- lsct command
 - ASMCMD utility, [10-14](#)

- lsdg command
 - ASMCMD utility, [10-61](#)
- lsdsk command
 - ASMCMD utility, [10-63](#)
- lsgrp command
 - ASMCMD utility, [10-96](#)
- lsod command
 - ASMCMD utility, [10-65](#)
- lsdf command
 - ASMCMD utility, [10-49](#)
- lsop command
 - ASMCMD utility, [10-15](#)
- lspwusr command
 - ASMCMD utility, [10-15](#)
- lstmpl command
 - ASMCMD utility, [10-88](#)
- lsusr command
 - ASMCMD utility, [10-96](#)

M

- managing a shared password file in a disk group
 - enabled by compatibility attributes, [4-58](#)
- managing aliases
 - Oracle ASM, [5-14](#)
- managing capacity in disk groups, [4-35](#)
- managing file access control
 - Oracle Enterprise Manager, [7-18](#)
- managing file access control for disk groups
 - Oracle Enterprise Manager, [7-18](#)
- managing Oracle ASM files
 - Oracle Enterprise Manager, [7-19](#)
- managing Oracle ASM templates
 - Oracle Enterprise Manager, [7-19](#)
- managing Oracle ASM templates page
 - Oracle Enterprise Manager, [7-19](#)
- managing templates
 - Oracle Automatic Storage Management, [5-19](#)
- managing users
 - Oracle ASM in Oracle Enterprise Manager, [7-3](#)
- md_backup command
 - ASMCMD utility, [10-66](#)
- md_restore command
 - ASMCMD utility, [10-67](#)
- memory management
 - Oracle Automatic Storage Management, [3-7](#)
- metadata block size
 - Oracle ACFS, [16-28](#), [16-53](#)
- metadata on Oracle ADVM accelerator volume
 - enabled by compatibility attributes, [4-59](#)
- migrating a database
 - to Oracle ASM, [3-37](#)
- migrating data
 - Oracle ASM with Recovery Manager, [8-1](#)
- migrating data to alternative storage
 - Oracle ASM, [8-9](#)
- migrating data to Oracle ASM
 - with Recovery Manager, [8-6](#)

- migrating to Oracle ASM
 - Oracle Enterprise Manager, [7-23](#)
- mirror
 - fast resync, [4-43](#), [7-8](#), [7-9](#)
- mirroring
 - failure groups, [4-38](#)
 - failure groups, and, [4-38](#)
 - Oracle ASM, [1-6](#)
 - with failure groups, [1-6](#)
- mkalias command
 - ASMCMD utility, [10-50](#)
- mkcc command
 - ASMCMD utility, [10-31](#)
- mkdgm command
 - ASMCMD utility, [10-69](#)
- mkdir command
 - ASMCMD utility, [10-51](#)
- MKDIR command rule, [16-98](#)
- mkfsg command
 - ASMCMD utility, [10-83](#)
- mkfs command
 - AIX, [16-44](#)
 - Linux environments, [16-28](#)
 - Solaris, [16-36](#)
- mkgrp command
 - ASMCMD utility, [10-97](#)
- mktmpl command
 - ASMCMD utility, [10-89](#)
- mkusr command
 - ASMCMD utility, [10-98](#)
- MMAPREAD command rule, [16-98](#)
- MMAPWRITE command rule, [16-98](#)
- modifying a file
 - disk regions, [4-26](#)
- monitoring Oracle ASM performance
 - Oracle Enterprise Manager, [7-22](#)
- monitoring Oracle Flex ASM
 - Oracle Enterprise Manager, [7-25](#)
- mount command
 - AIX, [16-46](#)
 - ASMCMD utility, [10-71](#)
 - Linux environments, [16-30](#)
 - Solaris, [16-38](#)
- mount model
 - Oracle ACFS, [11-5](#)
- mount registry
 - Oracle ACFS, [11-11](#)
- mounting and dismounting disk groups, [4-67](#)
- mounting disk groups
 - using force, [4-67](#)
- mounting Oracle ACFS file systems, [16-5](#)
- moving data between disk groups
 - Oracle ASM, [8-10](#)
- multipathing
 - with Oracle ASM, [2-3](#)
- multiple file creation form
 - Oracle ASM, [5-4](#)

N

namespace

Oracle ACFS, [11-5](#)

naming

disk groups, [4-10](#)

network file systems

storage resources, [2-1](#)

NFS

and Oracle ACFS, [17-3](#)

NOWAIT keyword

in REBALANCE clause, [4-19](#)

O

OCR in disk groups

enabled by compatibility attributes, [4-58](#)

offline command

ASMCMD utility, [10-72](#)

online command

ASMCMD utility, [10-73](#)

open files

and chgrp command, [10-92](#)

and chmod command, [10-93](#)

and chown command, [10-94](#)

OPENFILE command rule, [16-98](#)

operating system authentication

for Oracle Automatic Storage Management, [3-32](#)

Oracle ASM, [3-32](#)

SYSASM, [3-32](#)

ora.diskgroup.volume.acfs file system resource

Oracle ACFS, [11-19](#)

ora.drivers.acfs drivers resource

Oracle ACFS, [11-18](#)

Oracle ACFS

command-line tools for replication, [16-63](#)

command-line tools for tagging, [16-60](#)

See also Oracle Automatic Storage Management Cluster File System

Oracle ACFS acfsutil command-line tools

displaying help information, [16-2](#)

managing trace files, [16-2](#)

privileges to run, [16-2](#)

using, [16-2](#)

version information, [16-2](#)

Oracle ACFS automatic resize

enabled by compatibility attributes, [4-59](#)

Oracle ACFS compression

enabled by compatibility attributes, [4-59](#)

Oracle ACFS Compression

overview, [11-28](#)

Oracle ACFS Defragger

enabled by compatibility attributes, [4-60](#)

Oracle ACFS loopback support, [17-4](#)

Oracle ACFS NAS Maximum Availability eXtensions

Oracle ACFS, [11-44](#)

Oracle ACFS replication

configuring snapshot-based replication, [18-1](#)

Oracle ACFS replication (*continued*)

configuring ssh for use with snapshot-based replication, [18-1](#)

installing ssh and Cygwin on Windows, [18-5](#)

upgrading to snapshot-based replication, [18-12](#)

Oracle ACFS security

command rules, [16-98](#)

Oracle ACFS snapshot-based replication

enabled by compatibility attributes, [4-60](#)

Oracle ACFS sparse files

enabled by compatibility attributes, [4-59](#)

Oracle ACFS support for 4 K sectors

enabled by compatibility attributes, [4-59](#)

Oracle ADVM

See Oracle Automatic Storage Management Dynamic Volume Manager

Oracle ADVM accelerator volume

enabled by compatibility attributes, [4-59](#)

Oracle ADVM volume devices

creating, [16-5](#)

Oracle ADVM volumes

managing, [4-21](#)

Oracle ASM

instance administration, [3-17](#)

See also Oracle Automatic Storage Management

Oracle ASM Configuration Assistant

about, [9-1](#), [14-1](#)

addDisk command, [9-10](#)

commands to manage Oracle ACFS, [14-6](#)

commands to manage Oracle ASM, [9-5](#)

configureASM command, [9-6](#)

configureParameter command, [9-7](#)

configuring an Oracle ACFS, [14-3](#)

configuring an Oracle ADVM volume, [14-2](#)

configuring disk groups, [9-2](#)

configuring Oracle ACFS encryption, [14-4](#)

configuring Oracle ACFS for a database home, [14-5](#)

configuring Oracle ACFS security, [14-4](#)

createACFS command, [14-8](#)

createACFSSnapshot command, [14-7](#)

createDiskGroup command, [9-9](#)

createVolume command, [14-8](#)

creating a database home on Oracle ACFS, [14-5](#)

creating an Oracle ACFS file system, [14-3](#)

creating an Oracle ADVM volume, [14-2](#)

creating an Oracle ASM instance, [9-2](#)

creating disk groups, [9-2](#)

deleteACFSSnapshot command, [14-7](#)

deleteASM command, [9-8](#)

mounting or dismounting an Oracle ACFS, [14-3](#)

running in command-line mode, [9-4](#)

running in silent mode, [9-4](#)

starting, [9-1](#)

updating Oracle ASM instance parameters, [9-2](#)

Oracle ASM directories

administering, [5-1](#)

- Oracle ASM Dynamic Volume Manager
 - overview, [11-1](#)
- Oracle ASM Extended Disk Groups, [4-80](#)
- Oracle ASM file access control commands
 - ASMCMD utility, [10-90](#)
- Oracle ASM file group commands
 - ASMCMD utility, [10-78](#)
- Oracle ASM File Groups
 - properties, [4-76](#)
 - quota groups, [4-79](#)
- Oracle ASM files
 - administering, [5-1](#)
- Oracle ASM Filter Driver
 - about, [20-2](#)
 - and extended partition tables, [20-3](#)
 - and Oracle ASMLIB, [20-3](#)
 - configuring, [20-3](#)
 - configuring after installation, [20-5](#)
 - configuring during installation, [20-4](#)
 - deconfiguring, [20-12](#)
 - determining the state, [20-10](#)
 - managed by Oracle Enterprise Manager, [7-20](#)
 - managing labels, [20-11](#)
 - migrating, [20-7](#)
 - overview, [20-1](#)
 - post-installation configuration, [20-5](#)
 - pre-installation configuration, [20-3](#), [20-4](#)
 - updating AFD_DISKSTRING, [20-10](#)
 - updating ASM_DISKSTRING, [20-11](#)
- Oracle ASM Filter Driver management commands
 - ASMCMD utility, [10-32](#)
- Oracle ASM flex and extended disk groups
 - enabled by compatibility attributes, [4-58](#)
- Oracle ASM Flex Disk Groups
 - about, [4-73](#)
- Oracle ASM instances
 - administration, [3-17](#)
- Oracle ASM performance
 - Oracle Enterprise Manager, [7-22](#)
- Oracle ASM preferred read
 - failure groups, [4-46](#)
- Oracle ASM proxy instance
 - Oracle ACFS and Oracle ADVM, [19-1](#)
- Oracle ASM Quota Groups, [4-72](#)
- Oracle ASM templates
 - administering, [5-1](#)
- Oracle Automatic Storage Management
 - about audit trail, [3-36](#)
 - accessing files with the XML DB virtual folder, [5-17](#)
 - active session history sampling, [3-24](#)
 - adding aliases, [5-14](#)
 - administering, [3-1](#)
 - administering Oracle Flex ASM, [19-6](#)
 - advanced topics, [1](#)
 - alias file names, [5-6](#)
 - aliases, [5-14](#)
 - Oracle Automatic Storage Management (*continued*)
 - and Oracle Enterprise Manager, [7-1](#)
 - ASM_POWER_LIMIT initialization parameter, [3-11](#)
 - ASM_PREFERRED_READ_FAILURE_GROUPS initialization parameter, [3-12](#)
 - ASMSNMP, [9-2](#)
 - ASMSNMP user in Oracle Enterprise Manager, [7-3](#)
 - authentication, [3-27](#)
 - backing up initialization parameter files, [3-4](#)
 - command-line utility, [10-1](#)
 - components, [1-3](#)
 - concepts, [1-1](#)
 - configuring initialization parameters, [3-2](#)
 - connecting to an instance, [3-18](#)
 - copying initialization parameter files, [3-4](#)
 - creating a database example, [21-1](#)
 - creating a database in, [21-1](#)
 - creating a new directory, [5-16](#)
 - creating archive log files in, [21-4](#)
 - creating control file example, [21-3](#), [21-4](#)
 - creating control file in, [21-3](#)
 - creating database files in, [21-1](#)
 - creating files in the database, [5-9](#)
 - creating files using a default file location, [5-9](#)
 - creating redo logs in, [21-2](#)
 - creating SPFILE in a disk group, [3-5](#)
 - creating tablespaces in, [21-2](#)
 - data files with alias names, [5-7](#)
 - database initialization parameters, [3-16](#)
 - DB_CACHE_SIZE initialization parameter, [3-13](#)
 - diagnostic directory, [3-14](#)
 - DIAGNOSTIC_DEST initialization parameter, [3-14](#)
 - directories, [5-16](#)
 - discovering disks, [1-11](#), [4-33](#)
 - disk discovery, [4-33](#)
 - disk discovery string, [1-11](#)
 - disk failures in, [4-41](#)
 - disk groups, [1-6](#)
 - disk partitions for storage resources, [2-1](#)
 - disk regions, [4-26](#)
 - disks, [1-7](#)
 - dropping a directory, [5-17](#)
 - dropping aliases, [5-15](#)
 - dropping files, [5-15](#)
 - extents, [1-8](#)
 - failure groups, [4-40](#)
 - failure groups recommendations, [4-41](#)
 - file access control, [4-61](#)
 - file names, [5-3](#)
 - file names with template form, [5-7](#)
 - file templates, [1-11](#)
 - file types supported, [5-1](#)
 - files, [1-8](#)

Oracle Automatic Storage Management (*continued*)

- fully qualified file names, [5-5](#)
- home page on Oracle Enterprise Manager, [7-2](#)
- incomplete file name form, [5-8](#)
- incomplete file name template form, [5-8](#)
- initialization parameter files, [3-3](#)
- initialization parameters, [3-8](#)
- INSTANCE_TYPE initialization parameter, [3-14](#)
- instances, [1-3](#)
- integration with Oracle ACFS, [11-15](#)
- LARGE_POOL_SIZE initialization parameter, [3-14](#)
- logical volumes for storage resources, [2-1](#)
- LUNs for storage resources, [2-1](#)
- managing aliases, [5-14](#)
- memory management, [3-7](#)
- migrating a database to, [3-37](#)
- migrating data, [8-1](#), [8-3](#)
- migrating data to, [8-6](#)
- migrating data to alternative storage, [8-9](#)
- mirroring and failure groups, [1-6](#)
- moving data between disk groups, [8-10](#)
- moving initialization parameter files, [3-4](#)
- multipathing, [2-3](#)
- multiple file creation form, [5-4](#)
- network file systems for storage resources, [2-1](#)
- new features, *xxxv*
- one operating system group for all privileges, [3-28](#)
- operating system authentication for, [3-32](#)
- Oracle Database file mapping, [5-11](#)
- Oracle Flex ASM, [19-1](#)
- Oracle IO Server, [19-1](#)
- ORACLE_HOME, [3-18](#)
- ORACLE_SID, [3-18](#)
- overview, [1-1](#)
- password file authentication for, [3-32](#)
- preparing storage resources, [2-1](#)
- privileges, [3-28](#)
- PROCESSES initialization parameter, [3-15](#)
- recommendations for storage preparation, [2-4](#)
- Recovery Manager, and, [8-1](#), [8-3](#)
- referencing files in the database, [5-9](#)
- REMOTE_LOGIN_PASSWORDFILE initialization parameter, [3-15](#)
- renaming a directory, [5-16](#)
- renaming aliases, [5-14](#)
- rolling patches, [3-26](#)
- rolling upgrades, [3-25](#)
- search order for initialization parameter files, [3-3](#)
- separate operating system groups for privileges, [3-29](#)
- setting initialization parameters, [3-7](#)
- setting up Oracle Flex ASM, [19-6](#)
- SHARED_POOL_SIZE initialization parameter, [3-16](#)

Oracle Automatic Storage Management (*continued*)

- shutting down an instance, [3-22](#)
- shutting down an instance with OCR or voting files in a disk group, [3-22](#)
- shutting down an instance with Oracle ACFS, [3-22](#)
- single file creation form, [5-4](#)
- starting up an instance, [3-18](#)
- starting up an instance with incorrect SPFILE path, [3-20](#)
- storage considerations, [2-1](#)
- storage limits, [4-49](#)
- storage resources, [2-1](#)
- striping, [1-9](#)
- types of ASMCMD commands, [10-2](#)
- users page in Oracle Enterprise Manager, [7-3](#)
- using file names in SQL statements, [5-10](#)
- views, [6-1](#)
- XML DB virtual folder, [5-17](#)

Oracle Automatic Storage Management Cluster File System

- about, [11-5](#)
- about using, [16-2](#)
- accelerator volume, [11-43](#), [16-28](#)
- adding objects to security realms, [16-9](#)
- adding security administrators, [16-9](#)
- administering, [11-15](#)
- advanced topics, [1](#), [17-1](#)
- and external tables on Windows, [11-15](#)
- and NFS, [17-3](#)
- and Oracle Enterprise Manager, [13-1](#)
- and Oracle Restart, [17-6](#)
- auditing, [11-30](#)
- auditing security and encryption, [11-29](#)
- backup and restore, [11-14](#)
- basic steps to access a file system, [16-7](#)
- basic steps to create a file system, [16-5](#)
- basic steps to deregister, dismount, and disable a volume and file system, [16-23](#)
- basic steps to manage snapshots, [16-8](#)
- basic steps to manage with command-line tools, [16-2](#)
- basic steps to remove a file system and volume, [16-24](#)
- changing tag names, [16-16](#)
- command-line tools, [16-1](#)
- command-line tools for auditing, [16-81](#)
- command-line tools for encryption, [16-117](#)
- command-line tools for security, [16-86](#)
- command-line tools for snapshots, [16-125](#)
- creating file systems, [16-5](#)
- creating security realms, [16-9](#)
- creating security rule sets, [16-9](#)
- creating security rules, [16-9](#)
- creating snapshots with Oracle Enterprise Manager, [13-4](#)

- Oracle Automatic Storage Management Cluster File System (*continued*)
 - creating volumes and file systems with Oracle Enterprise Manager, [13-2](#)
 - database files with Oracle Exadata, [11-6](#)
 - deregistering file systems, [16-23](#)
 - determining the Oracle ADVM volume device name, [16-5](#)
 - disabling a volume, [16-24](#)
 - disk group compatibility, [16-5](#)
 - dismount or shut down, [11-22](#)
 - dismounting file systems, [16-24](#)
 - displaying encryption information, [16-14](#)
 - displaying security information, [16-9](#)
 - displaying tagging information, [16-16](#)
 - driver commands, [17-6](#)
 - driver model, [11-10](#)
 - drivers resource, [17-4](#)
 - enabling encryption, [16-14](#)
 - encrypting file systems, [16-14](#)
 - encrypting secure file systems, [16-9](#)
 - encryption, [11-25](#)
 - encryption keys, [11-25](#)
 - error handling, [17-3](#)
 - file access and security, [11-16](#)
 - file mapping, [12-6](#)
 - file systems on other nodes, [16-7](#)
 - grid infrastructure configuration, [11-17](#)
 - grid infrastructure installation, [11-17](#)
 - High Availability NFS for Oracle Grid Home, [11-44](#)
 - I/O failure console messages, [17-29](#)
 - important considerations with database files, [11-6](#)
 - important general considerations, [11-1](#)
 - individual file system resource, [17-5](#)
 - initializing encryption on file systems, [16-14](#)
 - initializing security on file systems, [16-9](#)
 - integration with Oracle ASM, [11-15](#)
 - limit on the number of snapshots, [11-11](#)
 - limits, [17-1](#)
 - managing encryption with Oracle Enterprise Manager, [13-16](#)
 - mount model, [11-5](#)
 - mount point and Oracle Database homes, [11-8](#)
 - mount registry, [11-11](#)
 - mounting file systems, [16-5](#)
 - namespace, [11-5](#)
 - new features, [xxxvii](#)
 - ora.diskgroup.volume.acfs file system resource, [11-19](#)
 - ora.drivers.acfs drivers resource, [11-18](#)
 - Oracle ACFS NAS Maximum Availability eXtensions, [11-44](#)
 - Oracle ADVM, [11-49](#)
 - Oracle ADVM volume devices, [16-5](#)
 - Oracle ASM proxy instance, [19-1](#)
 - Oracle Clusterware resource types, [11-18](#)
- Oracle Automatic Storage Management Cluster File System (*continued*)
 - Oracle Database homes, and, [11-8](#)
 - overview, [1-2](#), [11-1](#)
 - plug-in API, [17-8](#)
 - plugins, [11-43](#)
 - preparing security on file systems, [16-9](#)
 - privileges to run acfsutil command-line tools, [16-2](#)
 - read-only and read-write snapshots, [11-11](#)
 - read-write snapshots, [11-11](#)
 - registering file systems, [16-5](#)
 - registry resource, [17-5](#)
 - removing tag names, [16-16](#)
 - replicating file systems, [16-17](#)
 - replication, [11-36](#)
 - replication with auditing, encryption, and security, [11-41](#)
 - restrictions for storage resource, [2-1](#)
 - running acfsutil command-line tools on Windows, [16-2](#)
 - saving security information, [16-9](#)
 - securing file systems, [16-9](#)
 - security, [11-22](#)
 - setting encryption parameters, [16-14](#)
 - snapshots, [16-127](#), [16-131](#), [16-132](#)
 - space usage, [17-2](#)
 - specifying tag names for file systems, [16-16](#)
 - supported file types, [11-1](#)
 - system-created security realms, [16-97](#)
 - tagging, [11-39](#)
 - tagging API, [17-12](#)
 - tagging file systems, [16-16](#)
 - viewing and modifying snapshots with Oracle Enterprise Manager, [13-4](#)
 - viewing and modifying volumes and file systems with Oracle Enterprise Manager, [13-3](#)
 - views, [12-1](#)
- Oracle Automatic Storage Management Dynamic Volume Manager
 - accelerator volume, [11-43](#), [16-28](#)
 - default configuration, [15-2](#), [17-4](#)
 - limits, [17-4](#)
 - Oracle ACFS, [11-49](#)
 - overview, [11-10](#), [11-49](#)
- Oracle Cluster Registry (OCR)
 - in disk groups, [4-18](#)
- Oracle Clusterware resource types
 - Kernel Services Driver (OKS), [11-18](#)
 - Oracle ACFS, [11-18](#)
 - Oracle ADVM, [11-18](#)
- Oracle Database file mapping
 - with Oracle ACFS, [12-6](#)
 - with Oracle ASM, [5-11](#)
- Oracle Database homes
 - Oracle ACFS, and, [11-8](#)
- Oracle Enterprise Manager
 - accessing the Oracle ASM home page, [7-2](#), [13-1](#)

- Oracle Enterprise Manager (*continued*)
 - adding disks to a disk group, [7-7](#)
 - and Oracle ACFS, [13-1](#)
 - and Oracle ASM, [7-1](#)
 - backing up Oracle ASM files, [7-22](#)
 - bad block recovery in Oracle ASM disks, [7-23](#)
 - checking Oracle ASM disk group integrity, [7-22](#)
 - creating disk groups, [7-5](#)
 - creating Oracle ACFS file systems, [13-2](#)
 - creating Oracle ACFS snapshots, [13-4](#)
 - creating Oracle ACFS volumes, [13-2](#)
 - disk group compatibility, [7-13](#)
 - disk group file access control, [7-14](#)
 - disk group properties, [7-13](#)
 - disk group repair time, [7-14](#)
 - disk group smart scan compatibility, [7-14](#)
 - dismounting disk groups, [7-12](#)
 - dropping disk groups, [7-11](#)
 - dropping disks from disk groups, [7-9](#)
 - managing fast mirror resync, [7-8](#)
 - managing file access control for disk groups, [7-18](#)
 - managing Oracle ACFS encryption, [13-16](#)
 - managing Oracle ASM files, [7-19](#)
 - managing Oracle ASM Filter Driver, [7-20](#)
 - managing Oracle ASM templates, [7-19](#)
 - migrating to Oracle ASM, [7-23](#)
 - monitoring Oracle ASM disk group space usage, [7-12](#)
 - monitoring Oracle Flex ASM, [7-25](#)
 - mounting disk groups, [7-12](#)
 - Oracle ASM configuration page, [7-2](#)
 - Oracle ASM performance page, [7-22](#)
 - Oracle ASM users page, [7-3](#)
 - setting disk group compatibility, [7-5](#)
 - support workbench for Oracle ASM, [7-25](#)
 - viewing and modifying Oracle ACFS snapshots, [13-4](#)
 - viewing and modifying Oracle ACFS volumes and file systems, [13-3](#)
 - Oracle Flex ASM
 - administering, [19-6](#)
 - database clients, [19-1](#)
 - direct Oracle ASM clients, [19-1](#)
 - INSTANCE_TYPE initialization parameter, [19-4](#)
 - local Oracle ASM clients, [19-1](#)
 - managing, [19-1](#)
 - multiple Oracle ASM instances, [19-1](#)
 - Oracle ASM proxy instance, [19-1](#)
 - overview, [19-1](#)
 - relocating clients, [19-7](#)
 - setting up, [19-6](#)
 - Oracle Grid Infrastructure for a standalone server with Oracle ASM, [3-17](#)
 - Oracle Grid Infrastructure for independent servers
 - Oracle Restart, [3-17](#)
 - Oracle Grid Infrastructure for independent servers (*continued*)
 - See also* Oracle Grid Infrastructure for a standalone server
 - Oracle Home user on Windows
 - lower privileged user, [3-24](#)
 - Oracle IOServer
 - ASM_IO_PROCESSES initialization parameter, [3-11](#)
 - Oracle Flex ASM, [19-1](#)
 - Oracle Restart
 - and Oracle ACFS, [17-6](#)
 - See also* Oracle Grid Infrastructure for a standalone server
 - ORACLE_HOME
 - Oracle ASM, [3-18](#)
 - ORACLE_SID
 - Oracle ASM, [3-18](#)
 - orapwusr command
 - ASMCMD utility, [10-16](#)
 - OSASM group, [3-27](#), [3-29](#), [3-31](#), [3-32](#)
 - OSDBA group, [3-29](#)
 - OSOPER group, [3-29](#)
 - overview
 - ACFS
 - See* Oracle Automatic Storage Management Cluster File System
 - Oracle ACFS, [1-2](#), [11-1](#)
 - Oracle ACFS and Oracle ADVM, [11-1](#)
 - Oracle ADVM, [11-10](#)
 - Oracle ASM, [1-1](#)
 - OVERWRITE command rule, [16-98](#)
- ## P
-
- passwd command
 - ASMCMD utility, [10-98](#)
 - password file authentication
 - for Oracle Automatic Storage Management, [3-32](#)
 - Oracle ASM, [3-32](#)
 - performance and scalability improvements for ls and find
 - enabled by compatibility attributes, [4-59](#)
 - PHYS_META_REPLICATED disk group attribute, [4-7](#)
 - plug-in API
 - Oracle ACFS, [17-8](#)
 - Oracle ACFS pre-defined metric type, [17-8](#)
 - plug-in APIs
 - acfsfileid_lookup, [17-9](#)
 - acfsplugin_metrics, [17-9](#)
 - plugins
 - Oracle ACFS, [11-43](#)
 - Oracle ADVM, [11-43](#)
 - pre-defined metric type
 - Oracle ACFS plug-in API, [17-8](#)
 - preferred read
 - failure groups, [4-46](#)
 - preferred read failure groups
 - administering, [4-46](#)

- PREFERRED_READ.ENABLED
 - enabled by compatibility attributes, [4-58](#)
- PREFERRED_READ.ENABLED disk group attribute, [4-7](#)
- preparing security on file systems
 - Oracle ACFS, [16-9](#)
- privileges
 - ASMCMD utility, [10-9](#)
 - for one operating system group, [3-28](#)
 - Oracle ASM, [3-28](#)
 - separate operating system groups, [3-29](#)
 - SYSASM, [3-30–3-32](#)
 - SYSASM and Oracle Enterprise Manager, [7-2, 7-3](#)
 - SYSASM in Oracle Enterprise Manager, [7-5, 7-7](#)
 - SYSDBA with Oracle ASM, [3-31](#)
- PROCESSES initialization parameter
 - Oracle ASM instance, [3-15](#)
- pwdcopy command
 - ASMCMD utility, [10-17](#)
- pwcreate command
 - ASMCMD utility, [10-18](#)
- pwd command
 - ASMCMD utility, [10-51](#)
- pwddelete command
 - ASMCMD utility, [10-19](#)
- pwget command
 - ASMCMD utility, [10-19](#)
- pwmove command
 - ASMCMD utility, [10-20](#)
- pwset command
 - ASMCMD utility, [10-21](#)

Q

- quorum failure group, [4-18](#)

R

- READ command rule, [16-98](#)
- read errors
 - recovery, [4-42](#)
- read-only snapshots
 - enabled by compatibility attributes, [4-59](#)
- read-write snapshots
 - enabled by compatibility attributes, [4-59](#)
- READDIR command rule, [16-98](#)
- realms
 - system-created security, [16-97](#)
- rebal command
 - ASMCMD utility, [10-74](#)
- rebalance
 - and CONTENT.CHECK disk group attribute, [4-31](#)
 - and THIN_PROVISIONED disk group attribute, [4-31](#)
- REBALANCE NOWAIT clause, [4-19](#)
- REBALANCE WAIT clause, [4-19](#)

- rebalancing
 - BALANCE phase, [4-28](#)
 - COMPACT phase, [4-28](#)
 - disk groups, [1-13, 4-28](#)
 - phases of operation, [4-28](#)
 - PREPARE phase, [4-28](#)
 - RESTORE phase, [4-28](#)
 - RESYNC, RESILVER, and REBUILD operations, [4-28](#)
- rebalancing disk groups
 - tuning, [4-31](#)
- rebalancing process
 - ASM_POWER_LIMIT initialization parameter, [3-11](#)
- recommendations for storage preparation
 - for Oracle ASM, [2-4](#)
- recovery
 - read errors, [4-42](#)
 - write errors, [4-43](#)
- redo logs
 - creating in Oracle ASM, [21-2](#)
- redundancy
 - disk groups, [4-38](#)
- referencing files in the database
 - Oracle ASM, [5-9](#)
- regions
 - intelligent data placement, [4-26](#)
- registering
 - Oracle ACFS file system, [16-5](#)
- registry resource
 - Oracle ACFS, [17-5](#)
- remap command
 - ASMCMD utility, [10-75](#)
- remote instances
 - ASMCMD client cluster management commands, [10-30](#)
- REMOTE_LOGIN_PASSWORDFILE initialization parameter
 - Oracle ASM instance, [3-15](#)
- removing tag names
 - Oracle ACFS, [16-16](#)
- RENAME command rule, [16-98](#)
- renamedg tool
 - renaming disk groups, [4-70](#)
 - updating resources, [4-72](#)
- renaming a directory
 - Oracle ASM, [5-16](#)
- renaming aliases
 - Oracle ASM, [5-14](#)
- renaming disk groups
 - renamedg and updating resources, [4-72](#)
- renaming disks in a disk group, [4-24](#)
- replacing disks
 - altering disk groups, [4-24](#)
- replacing disks in a disk group, [4-24](#)
- replicated environments
 - disk group compatibility considerations, [4-60](#)
- replicating file systems

- replicating file systems (*continued*)
 - calculating storage requirements, [16-17](#)
 - Oracle ACFS, [16-17](#)
- replication
 - enabled by compatibility attributes, [4-59](#)
 - initiating, [16-69](#)
 - Oracle ACFS, [11-36](#)
 - storage requirements for Oracle ACFS file systems, [16-17](#)
 - using with Oracle ACFS auditing, encryption, and security, [11-41](#)
- replication status of a disk group
 - enabled by compatibility attributes, [4-58](#)
- resizing disks in disk groups, [4-27](#)
- restricting at startup, [3-21](#)
- resync
 - fast mirror, [4-43](#), [7-8](#), [7-9](#)
- reverting compatibility settings
 - disk groups, [4-55](#)
- rm command
 - ASMCMD utility, [10-52](#)
- rmalias command
 - ASMCMD utility, [10-53](#)
- rmcc command
 - ASMCMD utility, [10-32](#)
- RMDIR command rule, [16-98](#)
- rmgrp command
 - ASMCMD utility, [10-99](#)
- rmtmpl command
 - ASMCMD utility, [10-90](#)
- rmusr command
 - ASMCMD utility, [10-99](#)
- rolling patches
 - Oracle ASM, [3-26](#)
- rolling upgrades
 - Oracle Automatic Storage Management, [3-25](#)
- rpusr command
 - ASMCMD utility, [10-100](#)
- running
 - ASMCMD utility, [10-6](#)
 - ASMCMD utility in a database home, [10-6](#)
 - ASMCMD utility in an Oracle grid infrastructure, [10-6](#)

S

- saving security information
 - Oracle ACFS, [16-9](#)
- SCRUB_ASYNC_LIMIT
 - enabled by compatibility attributes, [4-58](#)
- SCRUB_ASYNC_LIMIT disk group attribute, [4-8](#)
- scrubbing disk groups, [4-32](#)
- sector size
 - disk groups, [4-15](#)
 - enabled by compatibility attributes, [4-58](#)
 - validations for disk groups, [4-15](#)
- SECTOR_SIZE disk group attribute, [4-8](#)
- securing file systems

- securing file systems (*continued*)
 - Oracle ACFS, [16-9](#)
- security
 - configuring with Oracle ASMCA, [14-4](#)
 - enabled by compatibility attributes, [4-59](#)
 - Oracle ACFS, [11-22](#)
 - Oracle ACFS and snapshots, [11-11](#)
 - Oracle ACFS command-line tools, [16-86](#)
 - system-created realms, [16-97](#)
- setattr command
 - ASMCMD utility, [10-75](#)
- setting encryption parameters
 - Oracle ACFS, [16-14](#)
- shared password file
 - for Oracle ASM instance, [3-34](#)
 - managing in a disk group, [3-34](#)
- SHARED_POOL_SIZE initialization parameter
 - Oracle ASM instance, [3-16](#)
- showclustermode command
 - ASMCMD utility, [10-22](#)
- showclusterstate command
 - ASMCMD utility, [10-22](#)
- showpatches command
 - ASMCMD utility, [10-22](#)
- showversion command
 - ASMCMD utility, [10-23](#)
- shutdown command
 - ASMCMD utility, [10-23](#)
- shutting down an Oracle ASM instance
 - about, [3-22](#)
 - with OCR or voting files in a disk group, [3-22](#)
 - with Oracle ACFS file system mounted, [3-22](#)
- single file creation form
 - Oracle ASM, [5-4](#)
- snapshot-based replication
 - Oracle ACFS, [11-36](#)
- snapshots
 - .ACFS directory, [11-11](#), [16-127](#)
 - converting format of an existing Oracle ACFS snapshot, [11-11](#)
 - creating from existing Oracle ACFS snapshot, [11-11](#)
 - creating Oracle ACFS read-only, [16-127](#)
 - creating Oracle ACFS read-write, [16-127](#)
 - duplicating, [16-131](#), [16-132](#)
 - limit on the number of, [11-11](#)
 - Oracle ACFS, [11-11](#), [16-127](#)
 - Oracle ACFS and encryption, [11-11](#)
 - Oracle ACFS and security, [11-11](#)
 - Oracle ACFS command-line tools, [16-125](#)
 - read-write in Oracle ACFS, [11-11](#)
 - snap directory, [11-11](#)
 - snaps directory, [16-127](#)
- Solaris
 - Oracle ACFS command-line tools, [16-33](#)
- space usage
 - Oracle ACFS, [17-2](#)

- spbackup command
 - ASMCMD utility, [10-24](#)
 - backing up an Oracle ASM SPFILE, [10-24](#)
- spcopy command
 - ASMCMD utility, [10-25](#)
 - copying an Oracle ASM SPFILE, [10-25](#)
- specifying tag names for file systems
 - Oracle ACFS, [16-16](#)
- spget command
 - ASMCMD utility, [10-27](#)
 - retrieving the location of an ASM SPFILE, [10-27](#)
- spmove command
 - ASMCMD utility, [10-27](#)
 - moving an Oracle ASM SPFILE, [10-27](#)
- spset command
 - ASMCMD utility, [10-28](#)
 - setting the location of an ASM SPFILE, [10-28](#)
- stamp command
 - ASMCMD utility, [10-76](#)
- stamplist command
 - ASMCMD utility, [10-77](#)
- standalone servers
 - Grid Infrastructure for a standalone server, [3-17](#)
- starting up an Oracle ASM instance
 - about, [3-18](#)
 - with incorrect SPFILE path, [3-20](#)
- startup
 - mounting disk groups, [3-21](#)
 - restricting disk groups, [3-21](#)
- startup command
 - ASMCMD utility, [10-29](#)
- storage considerations
 - for Oracle ASM, [2-1](#)
- storage limits
 - disk groups, [4-49](#)
 - Oracle ASM disks, [4-49](#)
 - Oracle Automatic Storage Management, [4-49](#)
- storage resources
 - disk partitions and Oracle ASM, [2-1](#)
 - for Oracle ASM, [2-1](#)
 - logical volumes and Oracle ASM, [2-1](#)
 - LUNs and Oracle ASM, [2-1](#)
 - network file systems and Oracle ASM, [2-1](#)
 - preparation for Oracle ASM, [2-1](#)
 - restrictions for Oracle ACFS, [2-1](#)
- STORAGE.TYPE disk group attribute, [4-8](#)
- storing data files and redo logs in Oracle ACFS file systems
 - enabled by compatibility attributes, [4-59](#)
- striping
 - files, [1-9](#)
 - Oracle ASM, [1-9](#)
- support for 1023 snapshots
 - enabled by compatibility attributes, [4-59](#)
- support for cloning, duplicating, remastering, and setting quotas for snapshots
 - enabled by compatibility attributes, [4-59](#)

- support workbench for Oracle ASM
 - Oracle Enterprise Manager, [7-25](#)
- supported file types
 - Oracle ACFS, [11-1](#)
- SYMLINK command rule, [16-98](#)
- SYSASM privilege
 - accessing the Oracle ASM home page in Oracle Enterprise Manager, [7-2](#)
 - Oracle Enterprise Manager, [7-3](#), [7-5](#), [7-7](#)
- SYSDBA privilege with Oracle ASM, [3-31](#)
- system-created realms
 - security, [16-97](#)

T

- tablespaces
 - creating in Oracle ASM, [21-2](#)
- tagging
 - enabled by compatibility attributes, [4-59](#)
 - Extended Attributes requirements, [11-39](#)
 - Oracle ACFS, [11-39](#)
 - requirements for Oracle ACFS, [11-39](#)
- tagging API
 - Oracle ACFS, [17-12](#)
- tagging error values
 - Oracle ACFS, [17-13](#)
- tagging file systems
 - Oracle ACFS, [16-16](#)
- tagging name specification
 - Oracle ACFS, [17-13](#)
- template management commands
 - ASMCMD utility, [10-87](#)
- templates
 - adding to a disk group, [5-22](#)
 - attributes, [5-20](#)
 - dropping, [5-23](#)
 - for specifying redundancy and striping, [5-23](#)
 - managing Oracle ASM, [5-19](#)
 - modifying an Oracle ASM, [5-23](#)
 - Oracle ASM, [1-11](#)
 - user-defined for Oracle ASM, [5-23](#)
- THIN_PROVISIONED disk group attribute, [4-9](#)
- tracking the replication status
 - for disk groups, [4-7](#)
- TRUNCATE command rule, [16-98](#)
- tuning rebalance operations, [4-31](#)

U

- umount command
 - AIX, [16-48](#)
 - ASMCMD utility, [10-78](#)
 - Linux environments, [16-32](#)
 - Solaris, [16-40](#)
- umountall command
 - AIX, [16-48](#)
 - Solaris, [16-40](#)

- undropping disks in disk groups, [4-28](#)
- unlimited file system expansion
 - acfsutil size command, [16-165](#)
 - enabled by compatibility attributes, [4-59](#)

V

- V\$ASM_ACFS_ENCRYPTION_INFO view, [12-1](#)
- V\$ASM_ACFS_SEC_ADMIN view, [12-2](#)
- V\$ASM_ACFS_SEC_CMDRULE view, [12-2](#)
- V\$ASM_ACFS_SEC_REALM view, [12-2](#)
- V\$ASM_ACFS_SEC_REALM_FILTER view, [12-2](#)
- V\$ASM_ACFS_SEC_REALM_GROUP view, [12-2](#)
- V\$ASM_ACFS_SEC_REALM_USER view, [12-2](#)
- V\$ASM_ACFS_SEC_RULE view, [12-2](#)
- V\$ASM_ACFS_SEC_RULESET view, [12-2](#)
- V\$ASM_ACFS_SEC_RULESET_RULE view, [12-2](#)
- V\$ASM_ACFS_SECURITY_INFO view, [12-2](#)
- V\$ASM_ACFSREPL view, [12-2](#)
- V\$ASM_ACFSREPLTAG view, [12-2](#)
- V\$ASM_ACFSSNAPSHOTS view, [12-2](#)
- V\$ASM_ACFSTAG view, [12-2](#)
- V\$ASM_ACFSVOLUMES view, [12-2](#)
- V\$ASM_ALIAS view, [6-1](#)
- V\$ASM_ATTRIBUTE view
 - enabled by compatibility attributes, [4-58](#)
- V\$ASM_AUDIT_CLEAN_EVENTS view, [6-1](#)
- V\$ASM_AUDIT_CLEANUP_JOBS view, [6-2](#)
- V\$ASM_AUDIT_CONFIG_PARAMS view, [6-2](#)
- V\$ASM_AUDIT_LAST_ARCH_TS view, [6-2](#)
- V\$ASM_CLIENT view, [6-2](#)
- V\$ASM_DISK view, [6-2](#)
- V\$ASM_DISK_IOSTAT view, [6-2](#)
- V\$ASM_DISK_STAT view, [6-2](#)
- V\$ASM_DISKGROUP view, [4-55](#), [6-2](#)
- V\$ASM_DISKGROUP_STAT view, [6-2](#)
- V\$ASM_ESTIMATE view, [6-2](#)
- V\$ASM_FILE view, [6-2](#)
- V\$ASM_FILEGROUP view, [6-2](#)
- V\$ASM_FILEGROUP_FILE view, [6-2](#)
- V\$ASM_FILEGROUP_PROPERTY view, [6-2](#)
- V\$ASM_FILESYSTEM view, [12-2](#)
- V\$ASM_OPERATION view, [6-2](#)
- V\$ASM_QUOTAGROUP view, [6-3](#)
- V\$ASM_TEMPLATE view, [6-3](#)
- V\$ASM_USER view, [6-3](#)
- V\$ASM_USERGROUP view, [6-3](#)

- V\$ASM_USERGROUP_MEMBER view, [6-3](#)
- V\$ASM_VOLUME view, [12-2](#)
- V\$ASM_VOLUME_STAT view, [12-3](#)
- variable size extents
 - enabled by compatibility attributes, [4-58](#)
- views
 - displaying Oracle ACFS information, [12-1](#)
 - displaying Oracle ASM information, [6-1](#)
- volcreate command
 - ASMCMD utility, [15-2](#)
- voldelete command
 - ASMCMD utility, [15-4](#)
- voldisable command
 - ASMCMD utility, [15-5](#)
- volenable command
 - ASMCMD utility, [15-6](#)
- volinfo command
 - ASMCMD utility, [15-6](#)
- volresize command
 - ASMCMD utility, [15-7](#)
- volset command
 - ASMCMD utility, [15-8](#)
- volstat command
 - ASMCMD utility, [15-9](#)
- volume management commands
 - ASMCMD utility, [15-1](#)
- volumes
 - enabled by compatibility attributes, [4-59](#)
- voting files in disk groups
 - enabled by compatibility attributes, [4-58](#)

W

- WAIT keyword
 - in REBALANCE clause, [4-19](#)
- wildcard characters
 - using quotes in ASMCMD commands, [10-10](#)
- Windows
 - Oracle ACFS command-line tools, [16-49](#)
- WRITE command rule, [16-98](#)
- write errors
 - recovery, [4-43](#)

X

- XML DB
 - virtual folder for Oracle ASM, [5-17](#)

