

Oracle ACFS 12c Release 2.

Introduction and Technical Overview

ORACLE WHITE PAPER | NOVEMBER 2017







Table of Contents

Introduction	1
Oracle ACFS High Level Architecture	2
Oracle ACFS Advanced Functionality	3
Oracle ACFS Security	3
Oracle ACFS Encryption	3
Oracle ACFS Compression	3
Oracle ACFS Auditing	3
Oracle ACFS Snapshots	4
Oracle ACFS Snapshot-based Replication	4
Oracle ACFS Tagging	5
Oracle ACFS Accelerator Volume	5
Oracle ACFS Defragging	6
Oracle ACFS Automatic Resize	6
Oracle ACFS NAS Maximum Availability eXtensions (MAX)	6
Oracle ACFS Sparse File Support	6
Oracle ACFS Performance	7
Oracle Platform Support	9
Oracle ACFS on Exadata Database Machine	9
Oracle ACFS Use Cases	10
Conclusion	11



Introduction

Oracle ASM Cluster File System (ACFS) is an industry standard, POSIX, X/OPEN and Windows compliant cluster file system. It supports multiple operating systems and server platforms including traditional server environments, the Oracle Exadata Database Machine and Oracle Database Appliance. It includes advanced features such as file system snapshots, replication, tagging, security, encryption, auditing, and highly available NFS and SMB services. Starting with Oracle Database 12c Release 2 features such as snapshot-based replication, compression, automatic resize, NAS Maximum Availability extensions, support for sparse files, metadata acceleration, among others, were introduced.

Oracle ACFS is integrated with Oracle Automatic Storage Management, extending its functionalities to support general-purpose files. Starting with Oracle Database 12c Release 1, ACFS also supports Oracle Database files. Oracle ACFS is designed to have direct access to Oracle ASM disk group storage for the shortest I/O path to deliver optimal performance. Through its underlying extent-based architecture, Oracle ACFS provides fast file access, high throughput and fast response time.

Oracle ACFS also includes Oracle ADVM, a dynamic cluster volume manager. Oracle ADVM provides volume management services and a standard device driver interface for system administrators to manage volumes across different nodes of a cluster with the same functionality on multiple platforms. Oracle ADVM volumes inherits key ASM features that includes striping, mirroring, rebalancing, intelligent data placement, preferred read, fast resync, even read, Flex ASM and many more.

Oracle ACFS leverages Oracle Clusterware for cluster membership state transitions and resource-based high availability. Oracle ACFS, ADVM, Oracle ASM and Oracle Clusterware are all bundled into the Oracle Grid Infrastructure (GI) allowing for integrated optimized management of databases resources, volumes and file systems.

Oracle ACFS provides a scalable, high performance and highly available solution that simplifies system and storage management through one management interface, one set of installation and configuration tools, one cluster framework and one vendor for support.

Oracle ACFS High Level Architecture

Oracle ACFS is a complete storage solution that offers performance, simple manageability and high availability. Oracle ACFS is integrated with Oracle ASM and Oracle Clusterware as shown in diagram 1 below. It communicates with Oracle ASM to store metadata locally which allows for direct access to ASM disk group storage for maximum performance. It also communicates with Oracle Clusterware to facilitate resource management for ACFS.

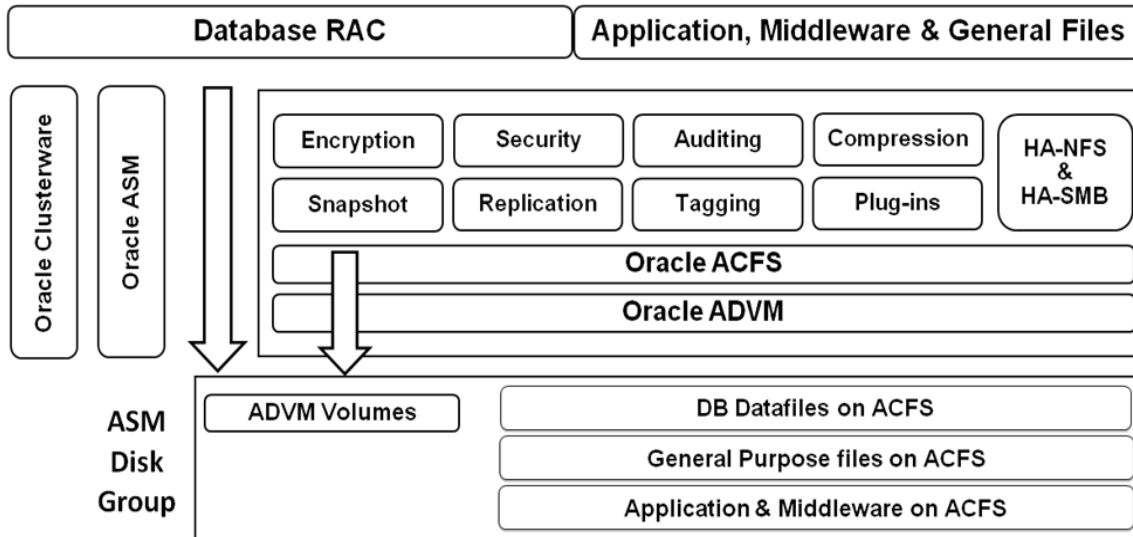


Diagram 1.

Upon creation of an Oracle ADVM volume a dynamic volume file is created for interface purposes of Oracle ACFS with Oracle ASM storage. During the creation of a file system, Oracle ACFS will refer to this file, upon configuration and mounting, the file system will inherit all Oracle ASM features associated with the related Oracle ADVM volume. Communication between Oracle ACFS and the Oracle ASM instance will now be performed through the Oracle ACFS driver. However, all I/O will go directly from Oracle ACFS to the Oracle ASM storage.

Oracle ACFS Clusterware support allows for Clusterware resources to take over operations such as cluster membership state transitions, driver loading, mount and unmount of file systems, enabling and disabling of volumes and providing High Availability for both Oracle ACFS and Oracle ADVM resources.

Oracle ACFS supports up to 64 mounted file systems for 32-bit systems, and up to 256 mounted file systems for 64-bit systems. File support in Oracle ACFS is for 1 trillion files in a file system, and there is no directory limit.

Oracle ASM disk group compatibility attributes set the storage limits for both Oracle ACFS and Oracle ASM, for further information please refer to the Oracle ASM Administrator's guide¹.

¹ [HTTPS://DOCS.ORACLE.COM/DATABASE/122/OSTMG/OSTMG.PDF](https://docs.oracle.com/database/122/OSTMG/OSTMG.PDF)



Oracle ACFS Advanced Functionality

Oracle ACFS Security

Oracle ACFS, similar to Oracle Database Vault, provides realm-based security. A realm, under Oracle ACFS, is a group of files or directories; only defined user or set of users can access this group. Each realm has a set of rules to enforce fine grain access control beyond what is provided initially by the operating system. Oracle ACFS Security data is logged for auditing and diagnostic purposes.

Oracle ACFS Security, provides protection from unauthorized access to the following elements:

- » Oracle ACFS Security Objects
 - » Security realms, rules and rules sets

- » Oracle ACFS Security Directory
 - » Log files and security metadata backup files.

- » Oracle ACFS realm-secured user files and directories

Oracle ACFS Encryption

Oracle ACFS file system data can be kept in an encrypted format through ACFS Encryption, allowing the user to encrypt individual files, directories or entire file systems. By offering encryption, Oracle ACFS users are protected against any unauthorized use, even in case of theft or data loss through disk failures. Oracle ACFS allows for co-existence of encrypted and non-encrypted files in the same Oracle ACFS file system. Applications will work unchanged and unaffected by the presence of file encryption.

Oracle ACFS encryption offers two different encryption keys, a file encryption key for files and a volume encryption key for file systems and encrypting file encryption keys. For key storing purposes, Oracle ACFS supports Oracle Cluster Registry (OCR) and starting with Oracle Database 12c Release 1 (12.1.0.2), Oracle Key Vault.

Oracle ACFS Compression

Oracle ACFS 12c Release 2 introduces Compression to its standard set of features. Oracle ACFS Compression is derived from Oracle's Database Advanced Compression solution. This solution comprises different sets of files such as RMAN files, data pump files, general purpose files and archive logs, this excludes redo logs, control files and flashback logs. Oracle ACFS Compression allows the user to reduce its storage footprint. Once Oracle ACFS compression is enabled on a file system, only files created afterwards are compressed, also, when disabling Oracle ACFS compression, compressed files are not automatically uncompressed.

Oracle ACFS Auditing

Oracle ACFS security and encryption can be audited via Oracle ACFS Auditing. All sources of events, in the case of security and encryption, are called audit sources, while all audit records are written in logs called audit trails. Audit trails are handled separately for every Oracle ACFS file system; on a cluster environment this include separation by individual nodes.

Oracle ACFS Auditing can be enabled at the realm level, this allows for auditing of authorizations and security violations. Oracle ACFS Auditing provides the entire underlying framework for the import of data into the Database Firewall and the Oracle Audit Vault by an audit vault collector. Oracle ACFS Auditing enables for separation of duties of the management and review of the audit to be enforced.

Oracle ACFS Snapshots

Oracle ACFS snapshotting functionality provides the ability to generate point in time copies of Oracle ACFS file systems. Oracle ACFS snapshots can either be read-only or read-write. Oracle ACFS snapshotting uses space-efficient copy-on-write functionality. In order to maintain a point-in-time view consistency of the file system, Oracle ACFS file extent values are preserved before they are modified in any way or even deleted.

Oracle ACFS Snapshots will be stored within the file system. If additional storage for files and snapshots is needed, Oracle ACFS allows for dynamic resizing of a file system. Upon creation, Oracle ACFS snapshots are ready and available to be used. Support for up to 1023 Oracle ACFS snapshots is available, snapshot of snapshots are supported; with full inheritance at any given point in the hierarchy. As long as the Oracle ACFS file system is mounted, snapshots will always be available.

Starting with Oracle ACFS 12c Release 2, further acfsutil operating system commands are introduced. The acfsutil snap duplicate commands are created to enhance replication. The acfsutil snap duplicate commands allows for create and apply operations.

The acfsutil snap duplicate create command, creates a duplicate snapshot of an already existing snapshot. The acfsutil snap duplicate apply command “applies” a duplicated snapshot to an existing file system or snapshot. This operation is an incremental apply, which means if it is the first time, the target file system or snapshot must be empty, then, after the first apply, all subsequent applies will only include deltas of the snapshot. Oracle ACFS snaputil commands allow now for snapshot remastering, renaming and quota setting.

Oracle ACFS Snapshot-based Replication

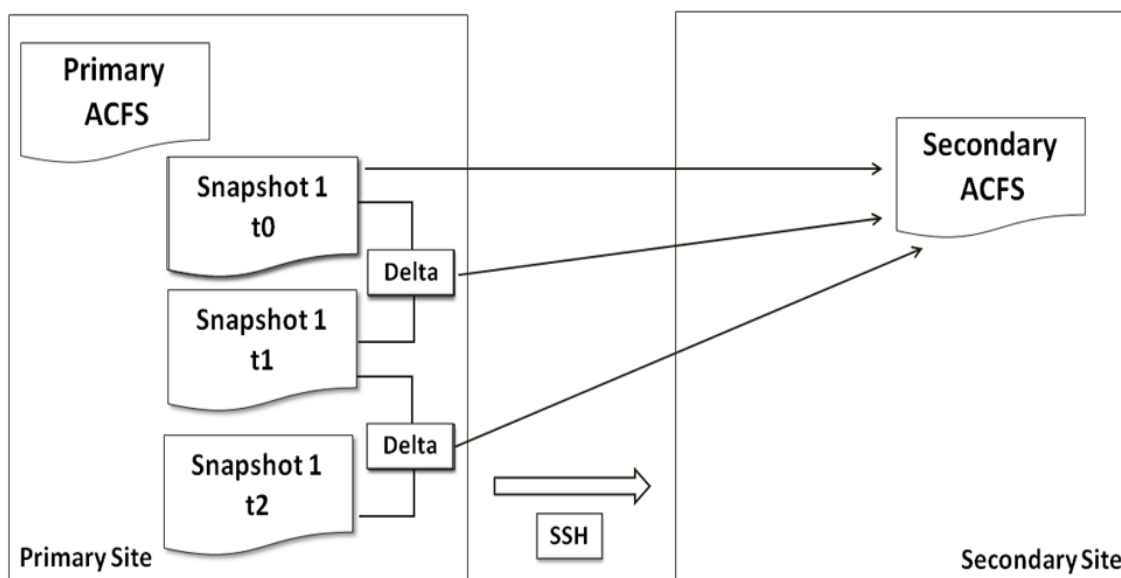



Diagram 2.



Oracle ACFS 12c Release 2 provides replication on a snapshot-based technology to provide maximum efficiency and flexibility in asynchronous replication. Oracle ACFS Snapshot-based Replication enables replication of an Oracle ACFS file system across the network from a primary to a standby site. Complemented with Oracle Data Guard, Oracle ACFS Snapshot-based Replication provides the customer with an end-to-end DR solution for all files.

Oracle ACFS Snapshot-based replication, works by transferring the differences between consecutive snapshots from the primary file system to the standby file system using the simple ssh protocol. Once an initial snapshot is replicated from the active to the standby file system, the changes of the successive snapshots will continue to be replicated. Oracle ACFS Snapshot-based replication carries a timestamp that can be queried for the purpose of synchronizing database transactions with file system data. Replication of Oracle Database files stored on Oracle ACFS file system is currently not supported with Oracle ACFS replication.

Oracle ACFS Replication on primary sites running AIX, Linux or Solaris works with standby sites running on any of the mentioned operating systems, but a primary site running on Windows is able to only replicate to a standby site also running on Windows. Oracle ACFS Replication allows for auditing, realm-based security and encryption to be enabled, thus securing the replicated standby file system with any of the policies in place in the primary file system.

Oracle ACFS Tagging

Oracle ACFS Tagging allows a user to associate one or more files as a group by assigning a unique 'tag name' attribute. Group operations can be performed based on tagged files that may span across different directories within an Oracle ACFS file system and within a single node or a cluster. By using Oracle ACFS Tagging, different groups of tagged files may be replicated as groups of related files, hence not having to replicate entire Oracle ACFS cluster file systems.

Oracle ACFS Accelerator Volume

Oracle ACFS Accelerator Volume was introduced in Oracle ACFS 12c Release 1 for future use. In Oracle ACFS 12c Release 2, this feature provides performance improvement for all metadata related operations. Since metadata is accessed more frequently than data, locating metadata in faster storage, often called an accelerator volume, will improve performance. ADVM region logging DRL files can also be stored on accelerator storage which optimizes IO performance for ADVM storage hosting Oracle ACFS file systems. An Oracle ACFS Accelerator Volume should use faster storage than the one used for the primary volume. The use of an accelerator volume is recommended for files with many extents; for instance, when there are Oracle ACFS Snapshots in use or for compressed ACFS file systems, the use of an Accelerator Volume is recommended. For further recommendations on guidelines for accelerator volumes size, please refer to the Oracle ASM Administrator's guide 1.

Oracle ACFS Defragging

Oracle ACFS is designed to eliminate the need for file system defragging. However, over time, a file system may need to be defragmented, specifically when storage is shared among database and snapshots or when the systems are under constant active workloads. Starting with Oracle ACFS 12c Release 2 a defragging feature is introduced, working in an automatic fashion in the background.

Oracle ACFS Automatic Resize

Starting on Oracle ACFS 12c Release 2, the `acfsutil size` command is introduced to provide an option for automatic file system resizing. With the `acfsutil size` command, the user specifies how much a file system should grow in case the available free space reaches certain threshold, also, the user can specify the maximum size a file system can reach, once the threshold is reached, the system will be automatically resized. This is an important and highly needed new feature introduced in this release.

Oracle ACFS NAS Maximum Availability eXtensions (MAX)

Oracle ACFS 12c Release 2 introduces a set of High Availability extensions for common NAS protocols, improving NFS and SMB exported file system availability. Oracle ACFS Maximum Availability eXtensions, while in use, guarantee that NFS or SMB will run in High Availability mode. This means that as long as there is one cluster node available, the NFS or SMB exports will be available. Oracle ACFS NAS Maximum Availability eXtensions integrate with NAS protocols and the Oracle ACFS stack, meaning ease of use without any additional infrastructure.

The High Availability Network File Storage (HANFS) feature enables highly available NFS servers to be configured using Oracle ACFS clusters, thus providing continuous service of NFS v2/v3/v4 exported paths. The HANFS cluster configurations may be built from your existing infrastructure or commodity servers and storage. Oracle ACFS now also supports HANFS NFS v4 with NFS Locks (limited to certain operating systems only – see documentation for details ²).

Oracle ACFS NAS Maximum Availability eXtensions offer support for SMB provided to interface with Active Directory Domains and Microsoft servers. Samba or Microsoft SMB must be in place in order to ensure Oracle ACFS HA-SMB. For further information please refer to Oracle ACFS NAS Maximum Availability eXtensions whitepaper ³.

Oracle ACFS Sparse File Support

With Oracle ACFS 12c Release 2, support for sparse files is introduced. Sparse files usually comprise of significant unused space, support for this files allows Oracle ACFS to leave a gap in the file, rather than allocating space on the storage for the unused space. This gap will then be “zero filled” in memory when read. This is key since it benefits NFS client write operations, NFS performance, disk utilization and storage saving for customers.

² <http://docs.oracle.com/database/122/OSTMG/understand-acfs-admin.htm#OSTMG95489>

³ <http://www.oracle.com/technetwork/database/database-technologies/cloud-storage/acfs/learnmore/acfs-nas-max-wp-3618364.html>

Oracle ACFS Performance

Oracle ACFS underlying architecture delivers optimal performance for general-purpose files and Oracle Database files. Oracle ACFS utilizes a direct I/O path from the database to the ASM disk group devices and hence bypasses the traditional OS VM page caching mechanism used for file system applications. This model enables ACFS-based databases to achieve performance metrics comparable to databases configured to use Oracle ASM directly. Comprehensive internal benchmarks using OLTP and DSS workloads have proven the performance of Oracle ACFS and Oracle ASM to be similar. Such benchmarks were ran using the following infrastructure:

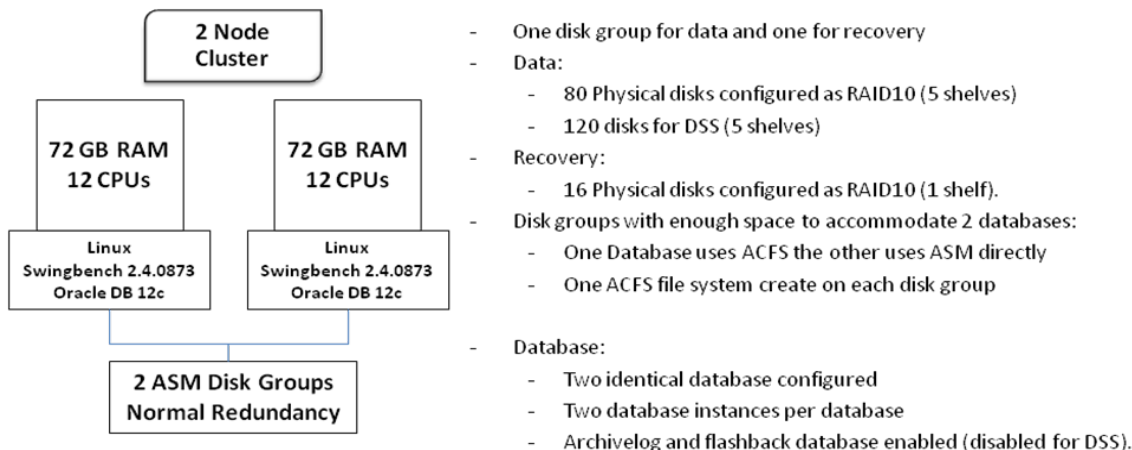



Diagram 3.



The benchmark configuration for Swingbench is comprised of two database schemas: OE SCHEMA (OLTP), with no partitioning, no compression, big file tablespace, all indexes, and sizing factor 6, meaning a 19 GB Tablespace, and the SH SCHEMA (DSS) created with range partitioning, no compression, big file tablespace, all indexes and sizing factor 41, meaning a 70 GB Tablespace.

The following charts shows performance based on an OLTP workload. As it can be seen, the throughput performance and transaction response time are almost identical regardless of whether the database use ACFS or ASM directly. The following comparison was made using ten processes and data is based on the average of 5 thirty minute test runs on Oracle ASM and Oracle ACFS.





The operating system and CPU loads show no significant difference between Oracle ACFS and ASM; neither did the DSS workload as shown in the following:



PERFORMANCE NUMBERS

	ASM	ACFS	% difference
SH Throughput (TX/min)	3.29	3.26	-1%
IO Throughput (Reads / sec)	3501	3312	-5%
IO Throughput (MB Read / sec)	724	712	-2%
CPU Usage (%)	16.5 %	16.7 %	+1%



Oracle Platform Support

Oracle ACFS is supported on multiple operating systems including Linux, Windows, Novell SLES, Solaris and AIX, for further information regarding kernel and Operating System specifics refer to My Oracle Support note 1369107.⁴ Oracle Engineered Systems where Oracle ACFS is supported include Oracle Database Appliance (ODA) and Oracle Exadata Database Machine.

Oracle ACFS on ODA

Oracle ACFS is the primary cluster file system to store database files and general purpose data in the Oracle Database Appliance. Since Oracle ACFS is built on top of Oracle ASM, the Oracle Database Appliance leverages the best of both products. Oracle Database Appliance adoption of Oracle ACFS provides with a wide array of functionalities requiring no additional effort or management.

File system management in the Oracle Database Appliance is done automatically by the integrated Appliance Manager, which means that all storage, including ACFS file systems and their underlying ASM diskgroups require no administration. By default, three file systems are created in the Oracle Database Appliance. These file systems are created in the already existing Oracle ASM disk groups, DATA, RECO and REDO. Oracle Database Appliance leverages Oracle ACFS snapshot capabilities, allowing for provision of test and development environments in the fastest and efficient way and performance equivalent to Oracle ASM.

Oracle ACFS on Exadata Database Machine

Oracle Exadata Database Machines run Oracle Linux and support Oracle ACFS for all database files and general purpose files starting with Oracle ACFS 12c Release 1 (12.1.0.2)

Oracle Database files supported in Oracle ACFS on Exadata Database Machines are limited to Oracle Database 10g Rel 2 (10.2.0.4 and 10.2.0.5), 11g (11.2.0.4 and higher) and Oracle Database 12c (12.1.0.1 and higher).

Oracle ACFS, as of now, does not use Oracle Smart Scan (predicate processing) and is not able to push database operations directly to storage leveraging Oracle Exadata storage. Oracle ACFS snapshot and tagging capabilities are supported for database files, however replication, security, encryption and audit capabilities are supported only for general-purpose files. For further information on Oracle ACFS on Exadata Database Machine refer to Oracle ASM Administrator's guide.¹

⁴ <https://support.oracle.com/epmos/faces/DocContentDisplay?id=1369107.1>

Oracle ACFS Use Cases

Oracle ACFS can be used in multiple scenarios that can leverage and benefit from its features and functionality. Oracle ACFS is ideal for Oracle RAC cluster environments, standalone deployments, Oracle middleware and Oracle Application environments and other vendor provided environments.

Oracle ACFS use cases range from managing Oracle Database files, general purpose files, shared database homes, administrative files, log files, trace and audit files, storing and managing business and unstructured data, delivering NFS/SMB exported file systems, leveraging NFS/SMB to off host backups, complementing Oracle Data Guard for Disaster Recovery solutions with Oracle ACFS Replication, leveraging snapshotting for quick provisioning test and development environments, among many other possible implementations.

Oracle Database files

- Leverages **standard**-based UI
- **Advanced** functionality

Shared homes

- For DB, Apps and middleware files
- HA and **simple** management

Business data

- Store business data on ACFS

HA/NFS server

- Leverage existing storage at **zero cost**

Backups

- **Simply** off host backups using ACFS NFS / CIFS

General purpose files

- OVM Repository, Golden Gate trail files, trace files, etc.

Conclusion

Oracle ACFS together with Oracle ASM offers a complete storage solution, a POSIX, X/OPEN and Windows compliant cluster file system and support of multiple operating systems and server platforms. It is used by customers in over 60 countries, used across all industries and leveraged by 40 of the top S&P 500 companies. Oracle ACFS simplifies management of Oracle Database files and general files and is an optimal solution for test and development environments.

Advanced functionality, such as replication, compression, tagging, security, auditing, resizing, metadata acceleration as well as NFS and SMB extensions add to the flexibility and value of Oracle ACFS. Based on those features, Oracle ACFS provides a complete storage solution delivering **performance**, **manageability** and best-in-class **availability**.

Performance	Manageability	Availability
<ul style="list-style-type: none">• Fast file access due to underlying extent-based architecture.• Throughput and response time go in line with ASM excellent performance numbers.	<ul style="list-style-type: none">• Industry standard POSIX & Windows Compliant Cluster File System• Easy management through command-line interfaces and graphical tools.• Oracle Clusterware integration allows automatic startup and mounting.	<ul style="list-style-type: none">• ACFS architecture allows quick recover from outages• Oracle Clusterware integration allows ACFS to leverage all cluster membership and high availability services.

Performance is ensured via fast file access, utilizing ACFS underlying extent-based architecture. Throughput and response time go in line with Oracle ASM excellent performance numbers.

Ease of management is provided by complying with industry standards, offering command-line interfaces as well as graphical tools for administration and its integration with Oracle Clusterware, allowing for automation of key tasks.

Best-in-class availability is ensured by its architecture and design that allows for quick recovery from outages and its integration with Oracle Clusterware allowing to leverage all cluster membership services.


Last but not least, Oracle ACFS is the best option for costumers looking for a complete and integrated storage solution designed to reduce total cost of ownership and simplification of system and storage management.



Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2017, Oracle and/or its affiliates. All rights reserved. This document is provided *for* information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle 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. 1117

Oracle ACFS 12c Release 2 – Introduction and Technical Overview
November 2017
Author: Oracle ACFS Product Management



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