ORACLE DATABASE 12c

An Oracle White Paper
January 2013

A Technical Overview of New Features for Automatic Storage Management in Oracle Database 12c

ORACLE

TABLE OF CONTENTS

Introduction

Automatic Storage Management (ASM) was first released with Oracle Database 10g. For Oracle databases, ASM provided a significant simplification for file system and volume management. In addition to enhancing storage automation, ASM improved file system scalability, performance, and database availability. These benefits hold for both single-instance Oracle databases as well as for Oracle Real Application Clusters (Oracle RAC) database environments. Oracle Database 12c introduces several new capabilities in ASM. This whitepaper presents these advances and is targeted at a technical audience mainly comprising:

- Database, system and storage administrators
- Architects
- Consultants
- System engineers
- Technical managers

ASM Overview

ASM is a feature of Oracle database providing an integrated cluster file system and volume manager at no additional cost. ASM lowers the Oracle database storage total cost of ownership and increases storage utilization while improving performance and availability as compared with traditional file system and volume management solutions. With ASM, less effort is required for managing your database storage environment.

ASM is easier to manage than conventional file systems, it optimizes the performance of your storage hardware, and is tightly integrated with the Oracle Database. Additionally, ASM eliminates the need for 3rd-party volume managers and file systems for managing the Oracle database files.

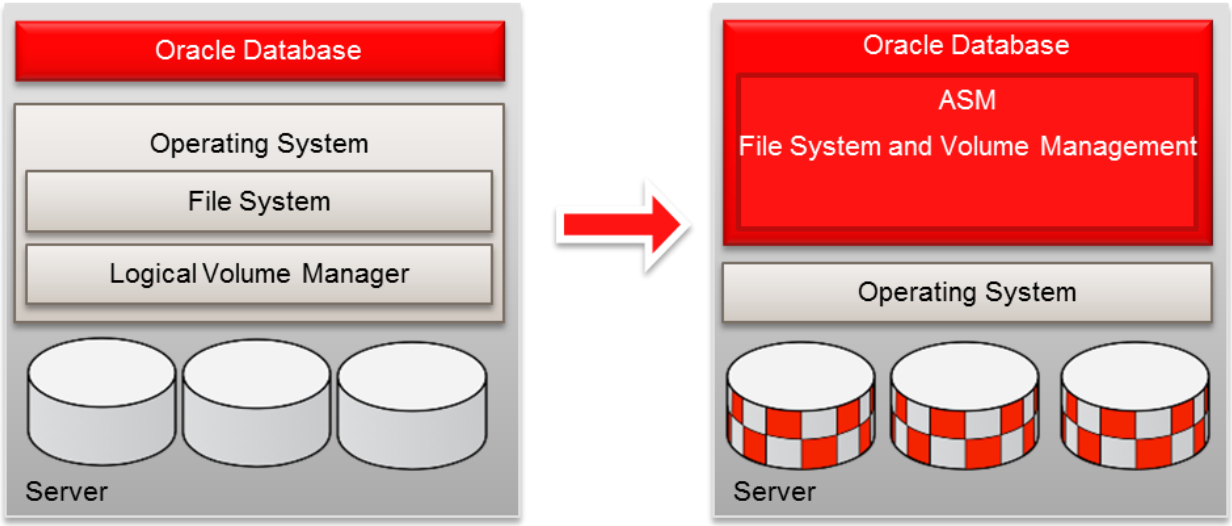

Figure 1

ASM brings significant key values to Oracle Database platforms. ASM improves manageability by simplifying storage provisioning, storage array migration, and storage consolidation. ASM provides flexible interfaces for management including the SQL*Plus, Oracle Enterprise Manager, and a UNIX-like command line interface called ASMCMD.

In addition to providing ease of management, ASM provides sustained best in class performance because of its innovative rebalancing feature. ASM distributes data evenly across all storage resources after storage configuration changes, providing an even distribution of IO and optimal performance. ASM scales to very large databases efficiently without compromising functionality or performance.

ASM is built to maximize database availability. ASM provides self-healing automatic mirror reconstruction and resynchronization, rolling upgrades and patching. ASM also supports dynamic and on-line storage reconfigurations in both single instance and Oracle RAC database configurations. ASM customers realize significant cost savings and achieve lower total cost of ownership because of features such as just-in-time provisioning, and clustered pool of storage making it ideal for database consolidation. ASM provides all of this without additional license or licensing fees.

In summation, ASM is a file system and volume manager optimized for Oracle database files providing:

- Simplified and automated storage management
- Increased storage utilization, uptime, and agility
- Delivering predictable performance and availability service levels

Total Storage Management for the Oracle Platform

With Oracle Database Release 11.2, Oracle added ASM Cluster File System (ACFS) to compliment Automatic Storage Management. ACFS provides the same level of storage management provided by ASM. ACFS simplifies and automates storage management functions, increases storage utilization, uptime and agility to deliver predictable performance and availability for file data stored outside an Oracle Database.

- ACFS includes: Automatic Storage Management Dynamic Volume Manager as a volume manager for Automatic Storage Management Cluster File System and 3rd party file systems.
- Automatic Storage Management Cluster File System provides advanced data services and security features for managing general purpose files.
- ASM, ACFS, and Oracle Clusterware are bundled as a complete package and is called *Oracle Grid Infrastructure*.

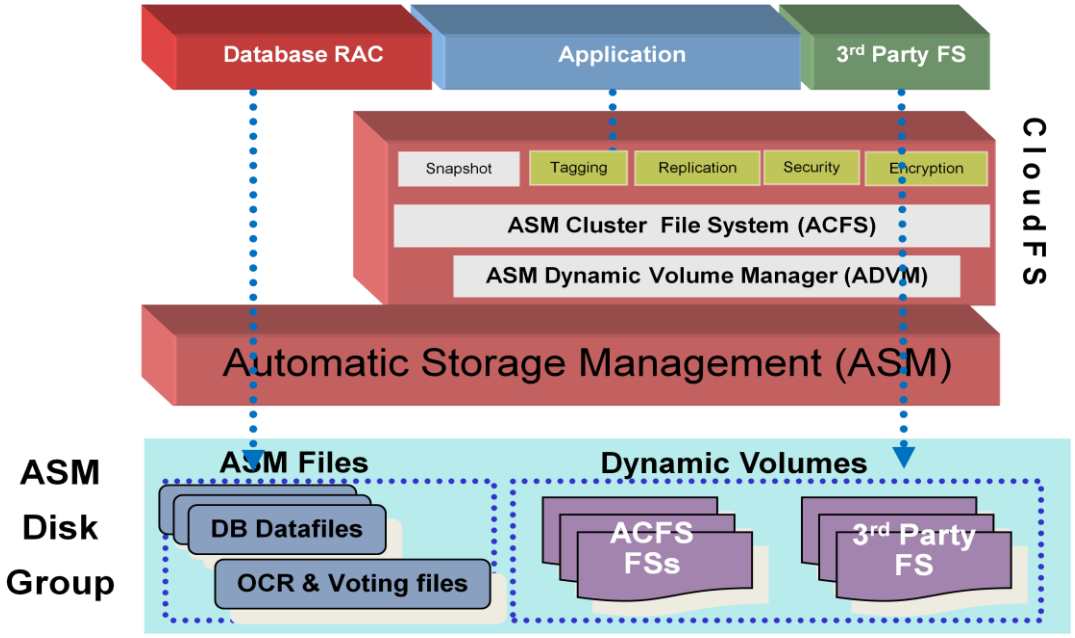

Figure 2

ASM Cluster File System, Automatic Storage Management and Oracle Clusterware are bundled as a single offering as *Oracle Grid Infrastructure*. Oracle Grid Infrastructure provides an integrated foundation for database and general purpose files as well as an infrastructure for clustered environments. The Oracle Grid Infrastructure streamlines management of volumes, file systems and cluster configurations, therefore eliminating the need for multiple 3rd party software layers that would add complexity and cost.

"With the ASM feature, DBAs won't have to worry about optimizing disk I/O. ASM eliminates hot spots by evenly distributing data providing the optimal bandwidth that an end user or an application needs" Arvind Gidwani,

Senior Manager IT, Qualcomm

"ASM solves the majority of the performance and manageability issues of using low-cost storage to run Oracle on Linux and allows you to build very large, high- performance systems" Grant McAlister

Principal Database Engineer Amazon.com

New ASM Features for Oracle Database 12c

Automatic Storage Management in the Cloud

When ASM was introduced in Oracle 10g, it was designed to address one goal really well; to improve the manageability of storage for Oracle databases. The next major evolution for ASM was in Oracle Release 11.2 with the introduction of ASM Cluster File System (ACFS). ACFS provided complete manageability for all customer data. In a continuing fashion, ASM in Oracle Database 12c addresses storage management needs for cloud computing in the enterprise. Cloud computing in the enterprise means applications and their supporting infrastructure are malleable within and across clusters and servers. New features in ASM for Oracle Database 12c enable storage management to seamlessly adapt to changing requirements and configurations. Furthermore, this evolution brings several new features greatly enhancing management for Oracle's engineered systems, such as Exadata and the Oracle Database Appliance.

Oracle Flex ASM

The most significant enhancement for ASM in Oracle Database 12c is a set of features collectively called Oracle Flex ASM. Oracle Flex ASM provides for critical capabilities required for cloud computing in enterprise environments. These environments typically deploy database clusters of varying sizes that not only have stringent performance and reliability requirements, but these environments must be able to rapidly adapt to changing workloads with minimal management overhead.

Oracle Flex ASM fundamentally changes the ASM cluster architecture. Before the introduction of Oracle Flex ASM in Oracle Database 12c, an ASM instance ran on every server in a cluster. These ASM instances communicated with other ASM instances on other servers in the cluster and collectively they presented shared Disk Groups to the database clients running in the cluster. This collection of ASM servers formed an ASM cluster. If an ASM instance were to fail, then all the database instances running on the same server as the failing ASM instance failed as well. The gray boxes in figure 3 represent ASM instances in a pre-12c environment.

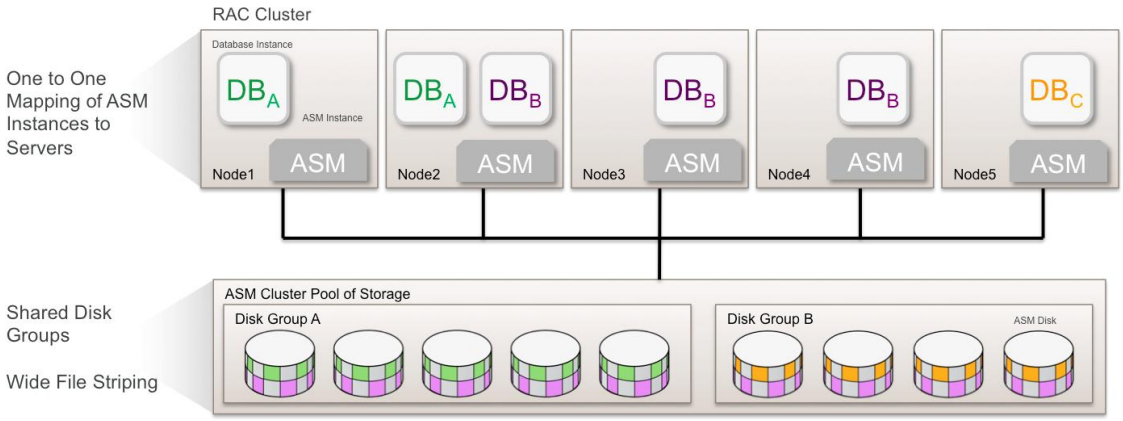

Figure 3

Oracle Flex ASM in Oracle Database 12c changes the architecture regarding the ASM cluster organization. In Oracle Database release 12c, only a smaller number of ASM instances need run on a subset of servers in a cluster. The number of ASM instances running is called the ASM cardinality. If a server fails that is running an ASM instance, Oracle Clusterware starts a replacement ASM instance on a different server to maintain the ASM cardinality. If an ASM instance fails for whatever reason, then

active Oracle 12c database instances that were relying on that ASM instance will reconnect to another surviving ASM instance on a different server. Furthermore, database instances are connection load balanced across the set of available ASM instances. The default ASM cardinality is 3, but that can be changed with a Clusterware command. These features are collectively called *Oracle Flex ASM*.

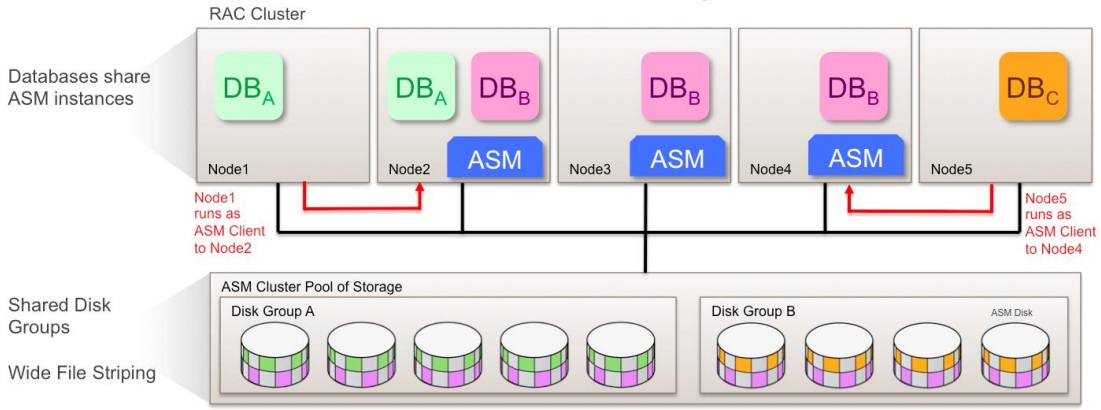

Figure 4: Flex ASM

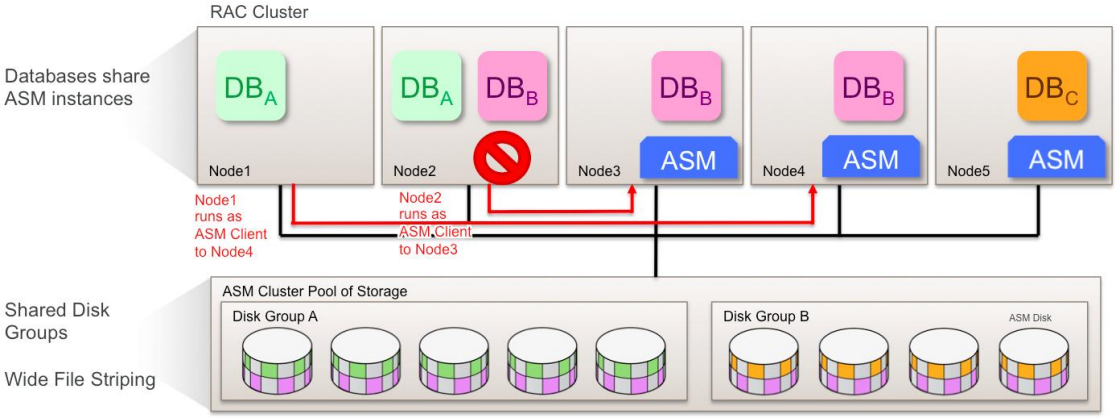

Figure 4a: ASM Instance failover

Figure 4 illustrates the ASM architecture with Oracle Flex ASM. There are a reduced number of ASM instances on selected servers in the cluster and Oracle Database 12c clients can connect across the network to ASM instances on different servers. Furthermore, Oracle Database 12c clients can failover to a surviving server with an ASM instance if a server with an ASM instance fails, all without disruption to the database client.

Dedicated ASM Network

Oracle Database 12c provides for the ability to dedicate a private network for ASM network traffic. The traffic on the ASM network is usually not overly significant and mostly metadata such as a particular file's extent map. If desired, the Oracle Clusterware private network can be shared with ASM and a dedicated ASM network is not required. The Oracle installer presents the DBA with the choice as to whether a dedicated network is required for ASM.

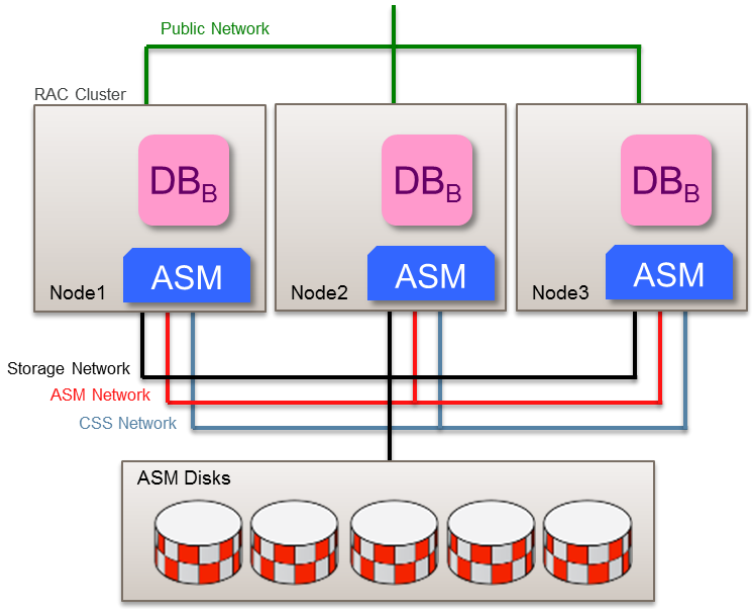

Figure 5

Remote Access

Before Oracle Database 12c, database clients use operating system authentication to log into an ASM instance. Historically, this approach worked well because the database instances and the ASM instance ran on the same server. With Oracle Database 12c, the database instances and ASM instances can now be on different servers, so ASM instances require a password file that is used to authenticate a database instance connecting to an ASM instance within the cluster. For most configurations a default ASM password file is created during installation and it is stored in an ASM Disk Group.

The ability of storing password files in a Disk Group is extended to Oracle Databases 12c clients. Having a common global password file in a cluster addresses common issues related to synchronizing multiple password files that had to be used previously.

Other Flex ASM Features

Other important features of Oracle Flex ASM include;

- The maximum number of ASM Disk Group is increased from 63 to 511.
- Oracle Flex ASM supports larger LUN sizes for Oracle Database 12c clients (increased to 32 PB).
- There is now a command for renaming an ASM Disk in a Disk Group.
- All the instances in an ASM cluster ensure they are running the same code release by validating the patch level across the cluster.

Deploying Oracle Flex ASM

For environments where ASM in Oracle Database 12c is deployed with databases running previous releases of Oracle, some new ASM features are not available to the older database clients. The critical capabilities unavailable are associated with pre-12c database clients not being able to access Oracle Flex ASM instances' running on remote servers from the server the database instance is running on. There are two models for how to deploy Flex ASM in such environments; In a *Greenfield Deployment*, Flex ASM and Oracle 12c databases are created in separate clusters. Over time applications and data are

relocated to the Oracle Flex ASM environment. This model offers the most separation between previous database releases and Oracle Database 12c. See figure 6.

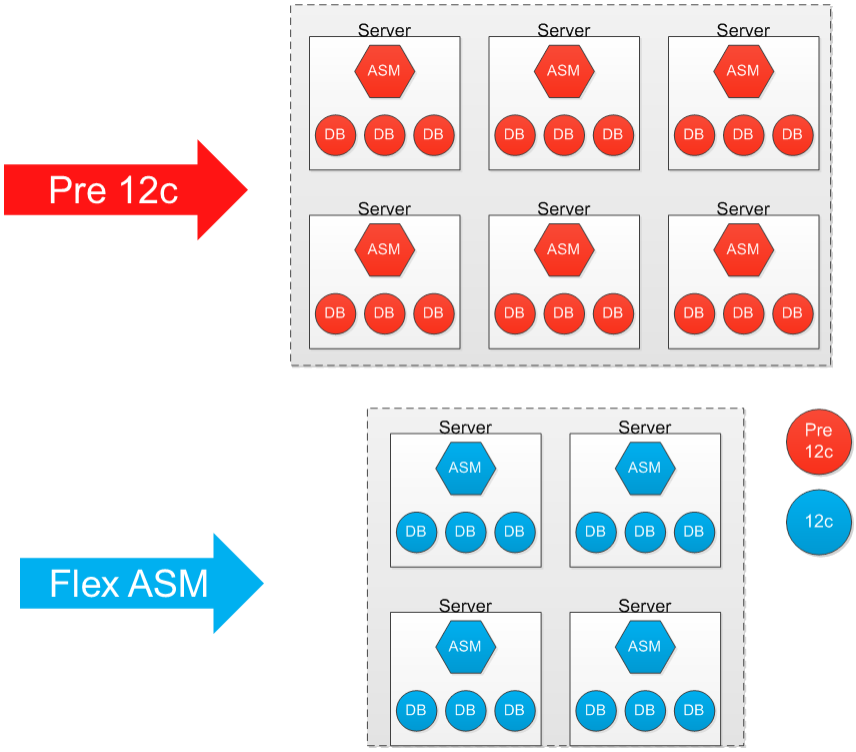

Figure 6

The second approach provides for a mixed environment of previous database releases and Oracle Database 12c clients in the same cluster. In this particular model, both Oracle release 12c and previous databases releases are operated in a cluster with an Oracle 12c ASM instance running on every server in the cluster. As before, ASM Disk Group compatibility attribute is use for managing the compatibility between database instances.

There are two ways to achieve this mixed model of operation. The first is to install the cluster in standard mode which assigns an ASM instance to every server. However, in standard mode, if an ASM instance fails, then Oracle 12c database instances using that ASM instance fail as well.

Another way to achieve the objective of an ASM instance on every server is to install Oracle Flex ASM and set the ASM cardinality to "all" which insures that there is an ASM instance on every server in a cluster. The advantage of this approach is that if an Oracle 12c database instance loses connectivity with an ASM instance, then the database connection will failover to another ASM instance on a different server. See figure 7.

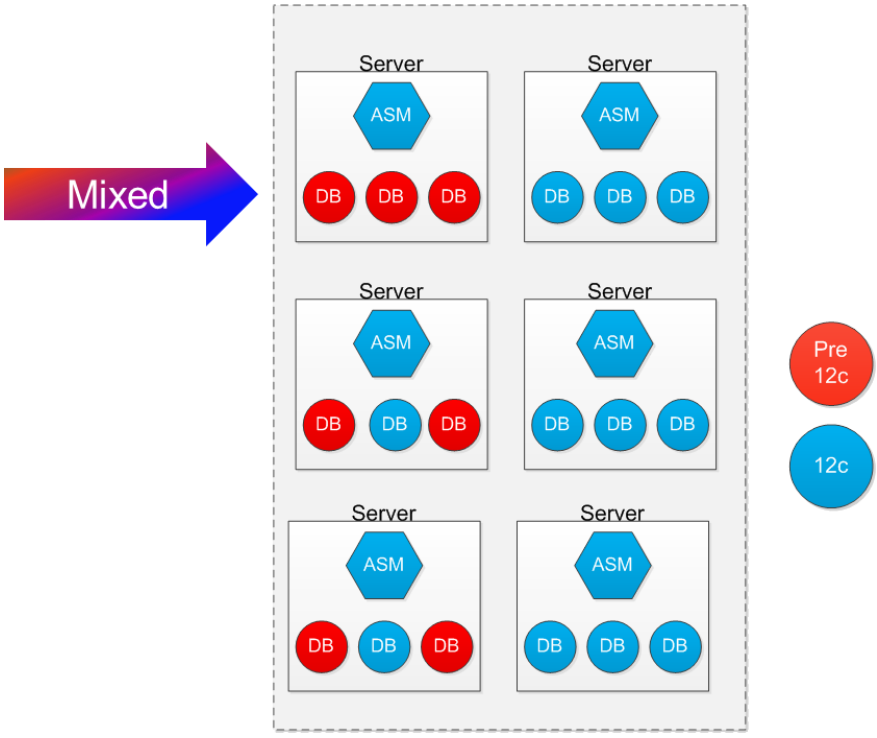

Figure 7

Managing Oracle Flex ASM

Oracle Flex ASM requires only minimal monitoring and ongoing management after initial configuration. The primary objective for administrators is that the ASM instances are up and running. This can be done using the Clusterware Oracle Clusterware **srvctl** command.

In Oracle Database 12c, there are no new instance parameters that are specific to Oracle Flex ASM. In addition, the default parameter settings have been adjusted to suit the Oracle Flex ASM architecture, making them sufficient to effectively support most situations.

ASM Features for Engineered Systems

There are several features in ASM for Oracle Database 12c facilitating the unique characteristics of Oracle's engineered systems, including Exadata and Oracle Database Appliance. While these features add value to Oracle's engineered systems, most of these features also provide similar capability for deployments of non-engineered systems. These features include;

Failure Group Repair Time

When an individual ASM Disk fails because of hardware fault, the failure is often terminal and the disk must be replaced. However, when all the ASM Disks in an ASM Failure Group fail simultaneously, it is unlikely that all the disks individually failed at the same time. Typically, it is more likely that some transient issue caused the failure associated with the ASM Failure Group. For example, all the disks in an ASM Failure Group in an Exadata environment would become unavailable when there is a transient failure of an Exadata storage cell. Because Failure Group outages are more likely to be transient in nature, and recovering redundancy through an ASM rebalance is far more expensive than replacing a single disk, it makes sense in these environments for the loss of a Failure Groups to have a larger repair time than that of an individual disk. A larger repair time value ensures that all the disks are not dropped automatically in the event of a short term and recoverable Failure Group outage. Consequently, ASM in Oracle Database 12c provides for a new Disk Group attribute called "failgroup_repair_time" that defaults to 24 hours.

ASM Disk Failure Handling Enhancements

Oracle 12c introduces a new feature that helps with the management associated with mirror synchronization in normal and high redundancy Disk Groups. This feature provides administrators with the ability to control the amount of resources dedicated to mirror re-synchronization. This is similar to the capability in previous ASM releases allowing administrators to control the use of resources dedicated to a Disk Group rebalance operations.

A related feature of resync in Oracle Database 12c ASM is that if a resync operation is interrupted and later restarted, the previously completed phases of the resync are skipped and processing is recommence at the beginning of the first remaining incomplete phase of the resync operation.

If an ASM Disk goes offline and cannot be repaired, administrators have to replace the disk. In prior versions of ASM, administrators had to drop the faulty disk and then add a new one back into the Disk Group. In the process, the entire Disk Group is rebalanced. This can be quite expensive with respect to moving data and time consuming. Oracle Database 12c ASM allows administrators to simply replace an offline disk using one fast and efficient operation. Using this feature, the replacement disk is populated with mirror copies of the ASM extents from other disks, and there is no need for any additional reorganization or rebalancing across the rest of the unaffected Disk Group.

In some situations, ASM rebalance operations are required to restore data redundancy within Disk Groups using normal or high ASM redundancy. For example, if a disk fails and no replacement disk is available, a rebalance is required to redistribute the data across the remaining disks in the Disk Group and restore redundancy. With Oracle Database 12c ASM, priority ordered rebalance is provided. This feature prioritizes on quickly restoring the redundancy of critical files first, such as control files and online redoes log files, to ensure that they are first protected against a secondary failure.

Since the first release of ASM, when data block is read, a series of checks are performed on data validating the block's logical consistency. If corruption is detected, the database automatically recovers by reading the mirror copy when normal and high redundancy Disk Groups are used. Extending this type of protection against hidden corruption for non-accessed data is a new feature for Oracle Database 12c. Under administrative control, ASM can proactively check for data corruption, even without any database client accessing the data. The value of proactive scrubbing is that without it, there is the possibility of multiple corruptions affecting all the copies of data that are infrequently accessed. Proactive scrubbing checks for, and where possible, repairs corruptions detected. Furthermore, this data checking can be triggered during rebalance operations or under execution of a scrubbing

command. On-demand scrubbing can be performed on a Disk Group, on individual files or individual disks.

ASM Data Allocation Enhancement

For normal and high redundancy ASM Disk Groups, the algorithm determining the placement of secondary extents uses an adjacency measure to determine the placement of mirror copies of data. In prior versions of ASM, the same algorithm and adjacency measure was used for all Disk Groups. Oracle Database 12c ASM provides administrators with the option to specify the content type associated with each ASM Disk Group. Three possible settings are allowed: *data*, *recovery* and *system*. Each content type setting modifies the adjacency measure used by the secondary extent placement algorithm.

The benefit is that the contents of Disk Groups with different content type settings are distributed across the available disks differently. This decreases the likelihood that a double-failure will result in data loss across normal redundancy Disk Groups with different content type settings. Likewise, a triple-failure is less likely to result in data loss for high redundancy disk groups with different content type settings.

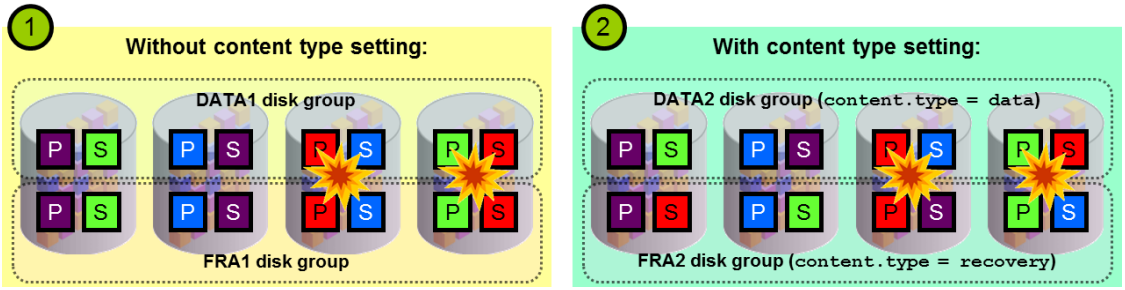

Figure 8

Extent Reading Selection Enhancement

In previous releases of ASM, the default behavior for ASM is to always read the primary copy of a mirrored extent unless a failure condition required otherwise. Alternatively, administrators could configure *preferred read failure groups* for extended cluster environments. For each I/O request presented to an ASM Disk Group with normal or high redundancy, there are one or more disks containing the data. The Even Read feature of Oracle Database 12c ASM, evenly distributes the selection of which copy of a block is read. With Even Read enabled, each read request is sent to the least loaded of the available disks.

Even Read is enabled by default on for Oracle Database 12c. Because Even Read is transparent to applications, users on I/O bound systems should notice a performance improvement after upgrading to version Oracle Database 12c. Note that Even Read does not replace the functionality provided by preferred read failure groups.

In an Exadata environment with Oracle Database 12c ASM, most extent relocations performed by a rebalance operation can be offloaded to Exadata Storage cell. With this capability, a single offload request can replace multiple read and write I/O requests. Offloading relocations avoids sending data to the ASM host, consequently improving rebalance performance.

Conclusion

ASM in Oracle Database 12c continues an evolution of superior storage management for the Oracle database that began with Oracle 10g. In Oracle 10g ASM was designed to address one goal really well; to improve the automation and manageability of storage for Oracle databases. The next phase for ASM's evolution was in Oracle Release 11.2 with the introduction of ACFS providing complete manageability for all customer data. ASM in Oracle 12c continues this tradition by addressing storage management requirements for cloud computing in the enterprise. New features in ASM for Oracle 12c enable storage management to seamlessly adapt to changing requirements and configurations. Part and parcel of this phase brings several new features enhancing management for Oracle's engineered systems, such as Exadata and the Oracle Database Appliance.

Oracle Whitepaper - A Technical Overview of New Features for Automatic Storage Management in Oracle Database 12c
January 2013
Author: Jim Williams
Contributing Authors:

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

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

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

Hardware and Software, Engineered to Work Together