ORACLE

A Technical Overview of New Features for Automatic Storage Management in Oracle Database 12c Release 2

ORACLE WHITE PAPER | FEBRUARY 2017

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Table of Contents

Introduction 1

Automatic Storage Management Overview and Background 1

Total Storage Management in Oracle 11g Release 2 2

New ASM Features Beginning in Oracle 12c Release 1 3

Oracle Flex ASM 3

New ASM Features in Oracle 12c Release 2 4

ASM IO Services 5

Database-Oriented Storage Management 6

File Groups 7

Quota Groups 8

ASM Database Clones 9

ASM Extended Diskgroups 9

Conclusion 11

Introduction

Automatic Storage Management (ASM) has been one of the most successful and widely adopted Oracle database features. ASM was introduced in Oracle Database 10g and through its many releases, ASM evolved to meet changing needs for both the Oracle database and the hardware environments in which the Oracle database operates. This paper presents the new ASM features in Oracle Database 12c Release 2 through a review of the evolution of ASM that existed in Oracle Database 10g to ASM in Oracle database 12c release 2. *If the reader's interest is primarily with the latest enhancements, then they can skip to the section titled "New Features in Oracle 12c Release 2" section.*

Automatic Storage Management Overview and Background

Automatic Storage Management (ASM) is a purpose-built file system and volume manager for the Oracle database. ASM was first released with Oracle Database 10g. For Oracle databases, ASM simplified both the file system and volume management for the Oracle database. In addition to simplifying storage management, ASM improved file system scalability, performance, and database availability. These benefits hold for both single-instance databases as well as for Oracle Real Application Cluster (Oracle RAC) databases.

With ASM, the customer does not need to use a third-party file system or a volume manager. Storage is provided to ASM for managing and ASM effectively organizes data in ASM Diskgroups.

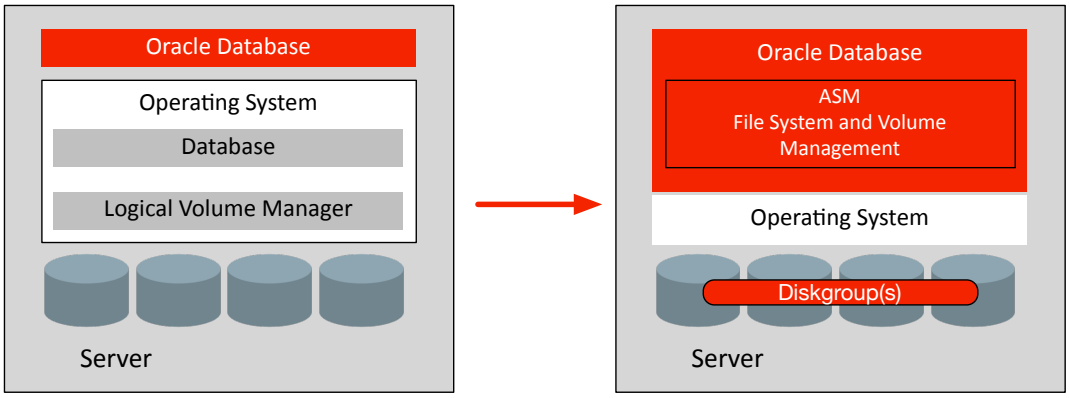

Figure 1. ASM replacing O/S Logical Volume Manager and File System

Because of its innovative rebalancing of data in Diskgroups, ASM maintains best in class performance by distributing data evenly across all storage resources whenever the physical storage configuration changes. This *rebalancing* feature provides an even distribution of IO and ensures optimal performance. Furthermore, ASM scales to very large configurations of both databases and storage, without compromising functionality or performance.

ASM is built to maximize database availability. For example, ASM provides self-healing automatic mirror reconstruction and resynchronization and rolling upgrades. ASM also supports dynamic and on-line storage reconfigurations. 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 licensing fees.

In summary, 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 in Oracle 11g Release 2

With the release of Oracle Database 11g Release 2, Oracle added the ASM Cluster File System (ACFS) to compliment ASM's file management. ACFS provides the same level of storage management provided by ASM. Specifically, ACFS simplifies and automates storage management functions, increases storage utilization, uptime and agility to deliver predictable performance and availability for conventional 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.
- 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 are called *Oracle Grid Infrastructure*

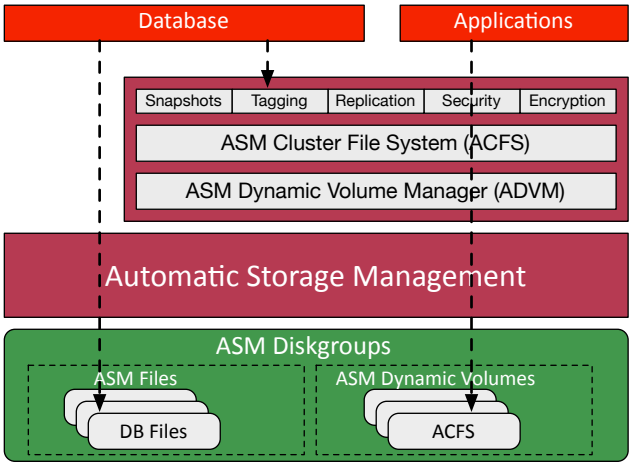

Figure 2. ASM System Layers

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 add complexity and cost.

New ASM Features Beginning in Oracle 12c Release 1

ASM in Oracle Database 12c Release 1 addressed storage management needs for environments with larger cluster configurations common to what has become called *private clouds*. ASM features introduced in Oracle database 12c Release 1 enhanced storage management by seamlessly adapting to frequently changing cluster configurations expected in cloud-like environments. Additionally, other features introduced enhancements 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 Release 1 is a set of features collectively called Oracle Flex ASM. Oracle Flex ASM provided for critical capabilities required for clustering in large enterprise environments. These environments typically deploy database clusters of varying sizes that not only have stringent performance and reliability requirements, but also must be able to rapidly adapt to changing workloads with minimal management.

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

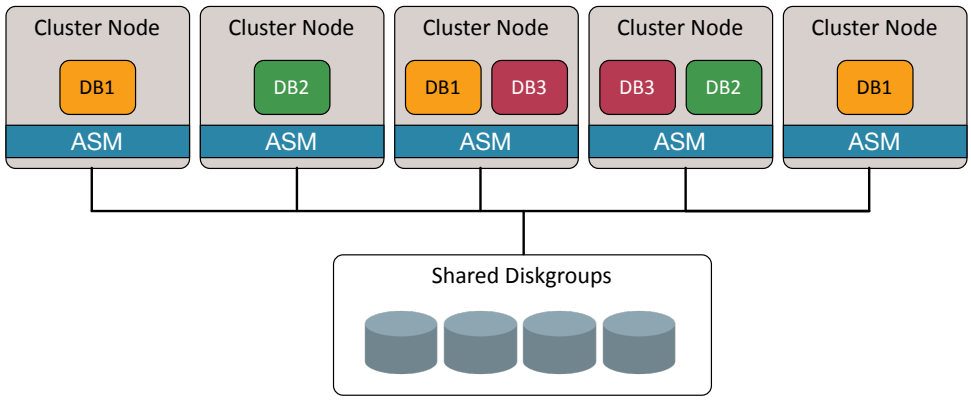

Figure 3. ASM Cluster

Oracle Flex ASM changed the architecture regarding the ASM cluster organization. In Oracle Database release 12c Release 1, a smaller number of ASM instances 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.

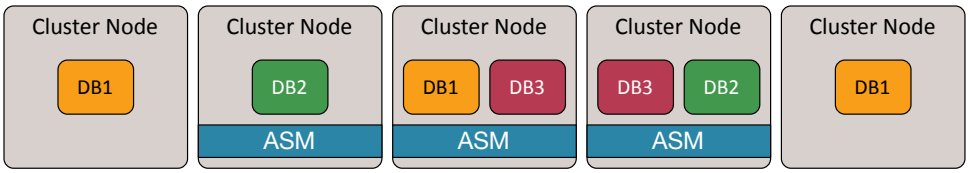

Figure 4. Flex ASM

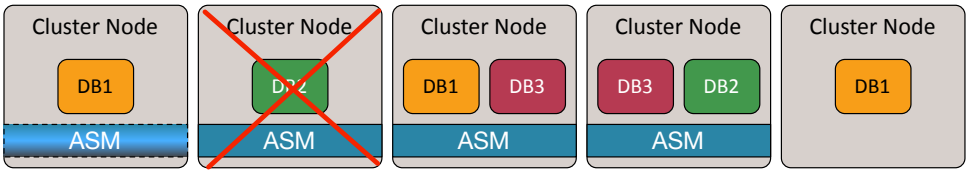

Figure 4a. Failing ASM Instance

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.

New ASM Features in Oracle 12c Release 2

The goal of ASM features introduced in Oracle Database 12c, Release 1 was to improve scaling and redundancy for large cluster environments. Enhanced scaling for clusters is further enhanced in Oracle Database 12c Release 2 with the introduction of *Cluster Domains*. A Cluster Domain provides a new architecture for clustering that centralizes a set of common core services utilized by multiple independent clusters, thus reducing overhead and streamlining manageability (see figure 5). It simplifies management of a large number of clusters and provides tools for enabling standardization. For example, provisioning, patching, diagnosticability, and storage management tasks are centralized and standardized. The key element in a Cluster Domain is the Domain Services Cluster (DSC). The DSC is a central point of management for a multi-clustered environment. The DSC provides a Management Repository that collects key statistics and diagnostic information useful for the individual Member Clusters and is the basis for the new Autonomous Health Framework.

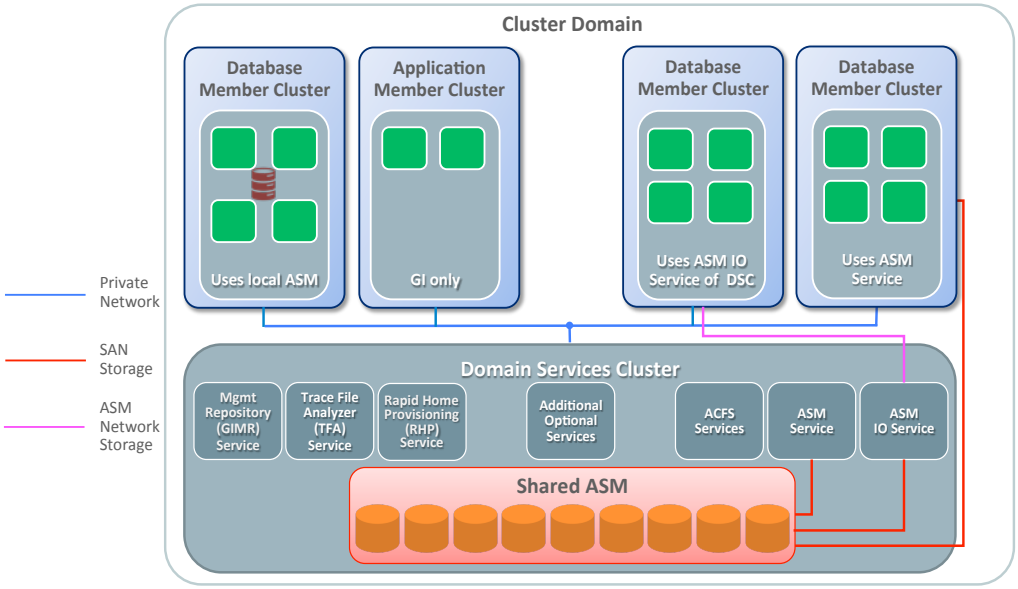

Figure 5. Domain Cluster

Simply stated, a *Cluster Domain* is collection of independent clusters, called *Member Clusters*, and a centralized *Domain Services Cluster* (DSC) providing centralized ASM and other management services. When a Cluster Domain is deployed, a DSC can eliminate the requirement for ASM instances to be running on the Member Clusters by utilizing ASM shared storage services in the DSC. However, while Member Clusters can use the ASM services in the DSC, as an alternative, Member Clusters, for storage isolation reasons, are able to host their own ASM environment. This is mode is shown in figure 5 with the left most Member Cluster.

ASM IO Services

ASM Services available in Oracle Database 12c release 2 enables Member Clusters to access data in a centralized ASM Diskgroup environment of the Domain Services Cluster (DSC). When ASM Services in a DSC are used, there are two connectivity options by which Members Clusters can access the shared Diskgroups residing in a DSC. The first is for the Member Cluster to have shared SAN attachment to the DSC storage as shown in figure 6a. The database instances in the Member Cluster rely on ASM instance(s) in the DSC for coordinating access through the SAN.

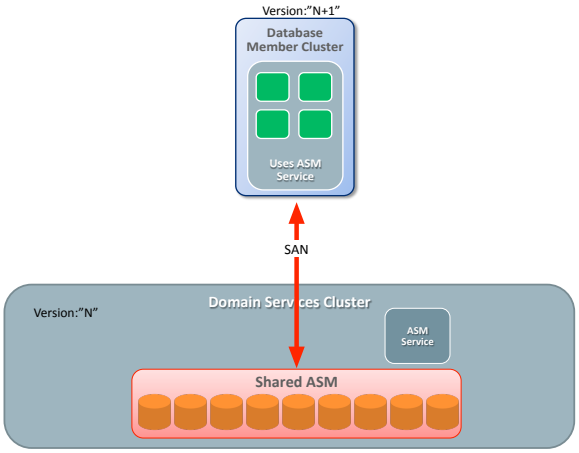

Figure6a. SAN Attached Storage

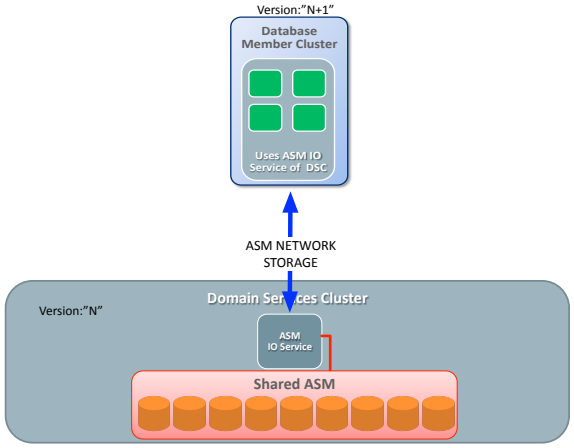

Figure 6b.ASM Network Attached Storage

The second way in which a Member Cluster can access a DSC's Diskgroup is though *ASM IO Services* using the *ASM Storage Network*. This is shown in figure 6b. With this mode, there is no physical SAN connection of storage between the Member Clusters and the DSC. Member Clusters access ASM data through a private network. The private network can be the same private network used by Grid Infrastructure or a separate dedicated *ASM private network*. This model of data access is useful for test and development configurations operating in Member Clusters for accessing production or test data that is shared through the DSC. The ASM Storage Network reduces the cost of storage with the use of network attachment rather than more expensive storage attachment, such as Fibre Channel switched networks.

Using a centralized ASM environment operating in a DSC allows Member Clusters to share ASM data and access shared Diskgroups. Customers can have a production database in the DSC, then clone that database using the new ASM database-cloning feature, and then use Rapid Home Provisioning feature for quickly creating a new test and development Member Cluster environment. Customers can quickly and easily create many test clusters without specialized or dedicated infrastructure. The key to these capabilities are the new ASM Services available in Oracle Database 12c Release 2.

Database-Oriented Storage Management

When introduced, the challenge for ASM was to address management complexity associated with storage used for large database environments. Before ASM, achieving optimal performance, database administrators had to determine the best storage configurations for various database objects. For example, a tablespace belonging to a particular table or database might best be placed on a particular file system providing the needed performance for the application. As databases grew is size, administrators had an ongoing challenge of making changes in the physical storage configuration to keep pace with the changing database workloads. ASM simplified this effort by allowing Diskgroups to contain all the database objects without regard to the underlying storage configuration. This meant that a particular Diskgroup contained a wide range of database files and even contained many different databases. This organization can be described as Diskgroup-oriented storage management. This is represented in figure 7 with many different databases and their files sharing a common Diskgroup. In this model there is no discrimination between different databases or even the needs of different databases.

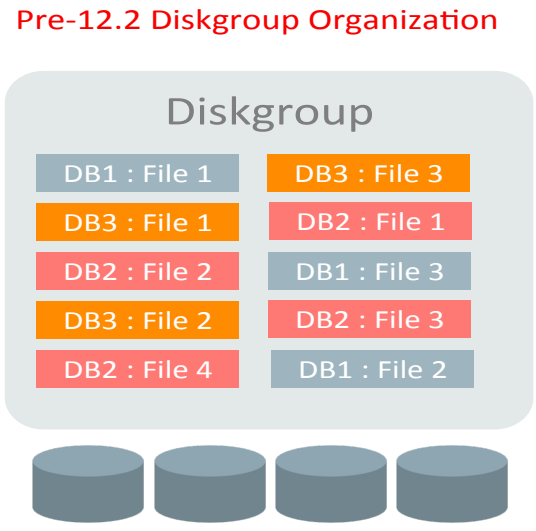

Figure 7: Disk-oriented Storage

Managing storage by allocating resources to a few Diskgroups did greatly improve the manageability for large databases and organizations deploying a large number of databases. It eliminated the piecemeal mode of the past

where administrators had to manage the relationship between database objects and storage with fine granularity. However, this course-grain management style limited the flexibility for effectively consolidating a large number of databases with different requirements into just a few simple Diskgroups using a single set of characteristics, such as redundancy. If was sometimes observed that some customers deployed many Diskgroups for containing databases that had different requirements. For example, production databases might be separated from test and development databases because the latter did not have the same performance or reliability requirements as the production databases. While separating databases into different Diskgroups provides finer granularity with respect to control, it works against the objective of reducing management overhead through consolidation.

File Groups

For these reasons, Oracle is introducing in Oracle Database 12c Release 2, the concept of "database-oriented storage management" with a new type of Diskgroups called Flex Diskgroups. With Flex Diskgroups, all files belonging to an individual database or with multitenet, a PDB are collectively identified with a new ASM object called a File Group. From here forward, for simplicity, we refer to non-multitent databases, and multitenant CDBs and PDBs as simply databases. A File Group logically contains the files associated with a single database. A File Group provides the means for referring to all the files that are part of a database residing within a single Diskgroup. A single database may have multiple File Groups residing in different Diskgroups. Command syntax referring to a File Group refers to all the files belonging to the File Group. Figure 8 below shows this logical grouping.

When a database is initially created, if an existing File Group already has a name that matches the name of the new database, then that existing File Group is used for recording the new file names that are created. However, if there is not an existing File Group for a database when a database is created, then a new File Group for the database is created.

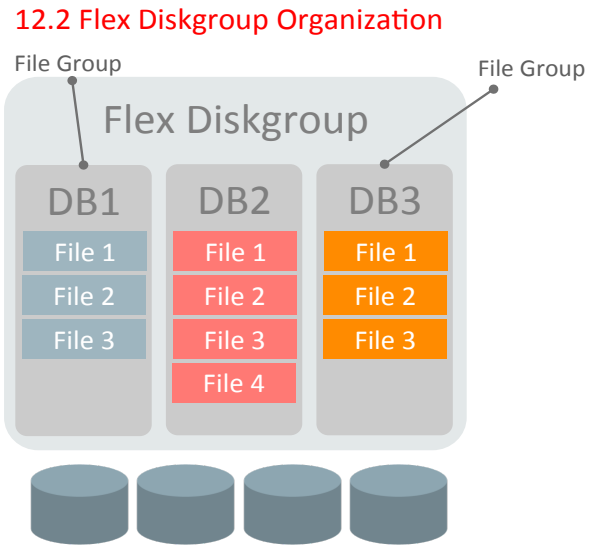

Figure 8: Database-oriented Storage Management

Different File Groups within a Flex Diskgroup can have different redundancies and the redundancies can be changed, as a situation requires. For example, a production database can use High Redundancy while a test database, in the same Diskgroup, can have Normal Redundancy. If desired, the redundancy of a database, i.e. the

File Group, can be changed online. When a File Group's redundancy is changed, ASM invokes an operation similar to an ASM rebalance causing the redundancy change in storage.
File Group names are unique within a Diskgroup. However, different Diskgroups can have File Groups with the same name. This provides for a database to span Diskgroups without creating a naming conflict.

Quota Groups

An important feature required for consolidating databases, from a storage management perspective, is storage quota management. Without the means of providing quota management, a single database can consume all the space is a Diskgroup. Flex Diskgroups offer a new feature called *Quota Groups*. A Quota Group is a logical container specifying the amount of Diskgroup space that one or more File Groups are permitted to consume. As an example, in figure 9 below, Quota Group A contains File Groups DB1 and DB2, whereas Quota Group B contains File Group DB3. The databases in Quota Group A are then limited by the specification of available space in that Quota Group.

Every Flex Diskgroup has a default Quota Group. If a File Group i.e. database, is not assigned a Quota Group, it is then assigned to the default Quota Group. Furthermore, the sum of space represented by all the Quota Groups may actually exceed the total physical space available. Consequently, Quota Groups represent a logical capacity limit of available space.

Changing Quota Groups requires ASM administrative privileges. An ASM administrator can create a set of Quota Groups in which subsequent databases are allocated. Quota Groups facilitate consolidating many databases into a single Flex Diskgroup by preventing any single database from consuming more than its fair share of storage and inhibiting the operation of the other databases.

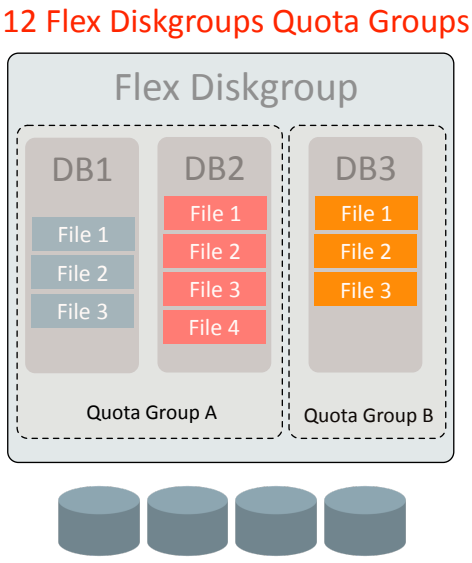

Figure 9: Quota Space Management

ASM Database Clones

A storage system's ability for rapidly replicating data for the purposes of deploying test and development systems is well understood and appreciated. Most storage-based replication technologies utilize either mirror splitting within storage arrays or snapshot replication in file systems. ASM with Flex Diskgroups now also provide the ability for creating near instantaneous provisioning of database copies. These database copies can be use for test and development, or when used with an Exadata system, provide a read-only master for an Exadata snapshot copies.

The advantage of ASM database clones, when compared with storage array-based replication, is that ASM database clones replicate complete databases rather than files or blocks of physical storage. Storage array or file system-based replication, in a database environment, requires coordination between database objects being replicated with the underlying technology doing the replication. With ASM database clones, the administrator does not need to understand the physical storage layout. This is another aspect of database-oriented storage management provided with ASM Flex Diskgroups.

ASM database cloning works by leveraging ASM redundancy. Previously, as a protection against data loss during hardware failure, ASM provided up to two additional redundant copies of a file's extents. Flex Diskgroups now can provide up to five redundant copies, in which one or more of the copies can be split off to provide a near instantaneous replica. The process involves two phases, the first phase is the replication phase and the second is the near instantaneous splitting of the copy providing a distinct replica that is independent of the original. When an ASM database clone is made, all the files associated with the database are split together and provide an independent database. Figure 10 represents the splitting of the files for DB3 providing a separate and independent database DB3a.

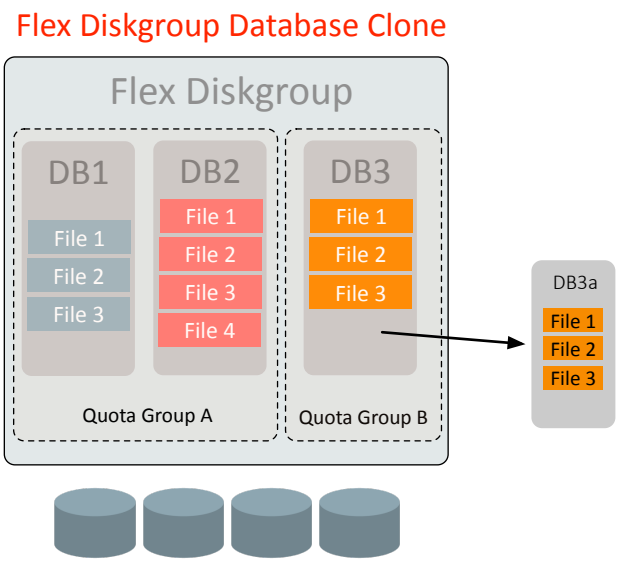

Figure 10: Database Clones

ASM Extended Diskgroups

With an enterprise's core operations at stake, it is vital that underlying storage for critical databases have the highest reliability and robustness. The design of ASM, from the beginning, has been to eliminate data loss after a hardware failure and ensure ongoing operation. For example, Normal and High-Redundancy Diskgroups provide data access in the face of storage failures. Early implementers of ASM saw an opportunity to extend an Oracle RAC cluster's availability by deploying RAC clusters across two closely located datacenters. This design used ASM mirroring

across the datacenters so that data availability is ensured in spite of a complete failure of a datacenter. Figure 11a represents the architecture of what became known as Extended RAC using ASM mirroring.

The benefit of the “Extended RAC” architecture is that ASM mirrors file extents across two different Failure Groups, with each File Group located in a different datacenter. If one datacenter fails, then all the file extents required for the RAC instance in the surviving datacenter, remains available in that datacenter.

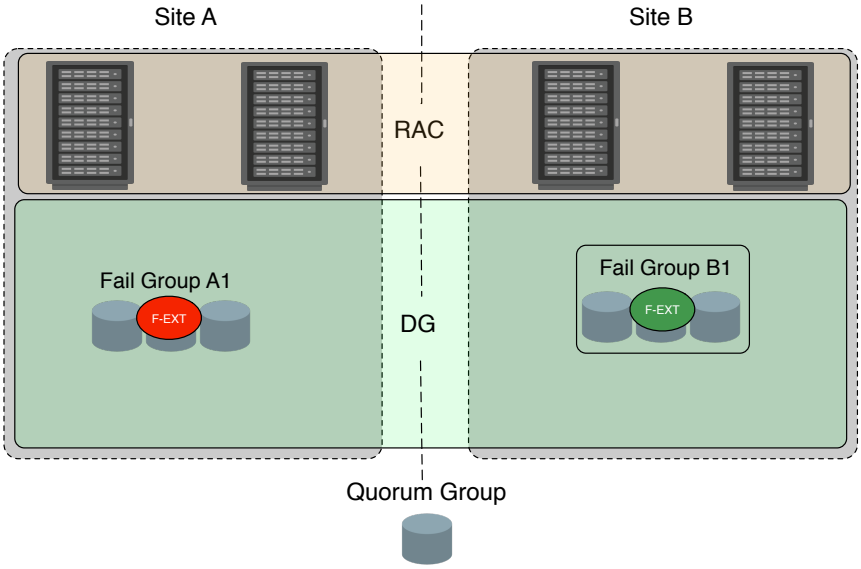

Figure 11a: ASM Extended Diskgroups

ASM in Oracle Database 12c Release 2 extends the utility of Extended RAC with a new feature called Extended Diskgroups. Extended Diskgroups eliminates the one limitation of Extended RAC, which was that a Diskgroup in an Extended RAC implementation could have at most two Failure Groups, each in a different datacenter. However, Extended Diskgroups, that are an extension of Flex Diskgroups, now enable multiple Failure Groups within a datacenter or site. This means that more than one copy of a file’s extent can exist, enabling mirroring within a datacenter, as well as across datacenters. In Figure 11b, the file extent high-lighted is replicated across Failure Groups within a datacenter as well as across the two datacenters. Additionally, Extended Diskgroups now support three datacenters and allow for the use of ASM High-Redundancy. Finally, with Extended Diskgroups, an Extended RAC configuration can be deployed on Exadata within the limitation of Infiniband.

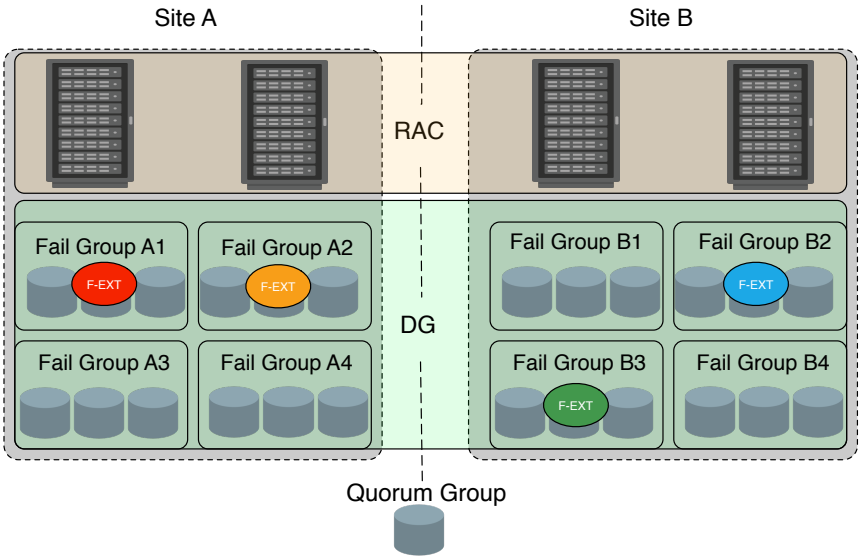

Figure 11b: ASM Extended Diskgroups

Conclusion

ASM in Oracle Database 12c Release 2 continues an evolution of superior storage management for the Oracle database that began with Oracle 10g. In Oracle 10g ASM addressed 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 Release 1 continued this tradition of evolution by addressing storage management requirements for larger cluster configurations including private cloud deployments in an enterprise. New features in ASM for Oracle 12c Release 2 further the tradition by supporting Domain Clusters. Domain Clusters and the Domain Services Cluster more broadly enable loosely joined and managed configurations that further reduces management overhead, yet provides for improved consolidation. The key to improving storage consolidation is Flex Diskgroups that enable database-oriented storage management that empower database management organizations and simplifies tasks of managing enterprise databases.

ORACLE

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

White Paper Title
February 2017
Author: [OPTIONAL]
Contributing Authors: [OPTIONAL]

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