

Automatic Storage Management Technical Overview

An Oracle White Paper
November 2003

Automatic Storage Management Technical Overview

Executive Overview	3
The DBA's Storage Management Challenge.....	4
Oracle 10G Provides a Solution –Automatic Storage Management	4
Removing Complexity	5
Distribute I/O to Optimize Performance.....	6
Intelligent Mirroring	6
How Automatic Storage Management Works	7
User Tunable/Configurable Parameters within ASM.....	7
Integrated with Oracle 10G components and features	8
Integration with Storage Vendors	8
Conclusion	8

Automatic Storage Management Technical Overview

EXECUTIVE OVERVIEW

Automatic Storage Management (ASM) is a new feature of Oracle Database 10g which provides a vertical integration of the file system and volume manager specifically built for the Oracle database files. ASM distributes I/O load across all available resource to optimize performance while removing the need for manual I/O tuning (spreading out the database files avoids hotspots). ASM helps DBAs manage a dynamic database environment by allowing them to grow the database size without having to shutdown the database to adjust the storage allocation.

Automatic Storage Management allows the DBA to define a pool of storage (called a disk groups) and then the Oracle kernel manages the file naming and placement of the database files on that pool of storage. The DBA can change the storage allocation (adding or removing disks) with new SQL commands (create diskgroup, alter diskgroup and drop diskgroup). One can also manage the disk groups by using Enterprise Manager (EM) and the Database Configuration Assistant (DBCA).

The Oracle Database 10g provides DBAs a simplified management interface for storage resources. Automatic Storage Management eliminates the need for manual I/O performance tuning. It virtualizes storage to a set of disk groups and provides redundancy options to enable a high level of protection. ASM facilitates non-intrusive storage configuration changes with automatic rebalancing. It spreads database files across all available storage to optimize performance and resource utilization. It is a capability that saves DBAs time by automating manual storage and thereby increasing their ability to manage larger databases and more of them with increased efficiency.

THE DBA'S STORAGE MANAGEMENT CHALLENGE

As the demand for database uptime continues to increase, so does the rate of growth in the average database size. These trends result in what some experts are calling “the management gap”. The fact that disk density and the database size keep doubling every few years is not being matched by doubling DBA staff in the same time period. Therefore, DBAs are being asked to manage more databases and larger databases with the same staff. They need tools that enable them to work more efficiently and manage more with less resource.

When it comes to managing how the database is laid out on storage, DBAs often face organizational challenges. DBAs often have to interact with Systems Administrators who in turn are working with Storage Administrators who might depend on the Network Administrator who is in charge of the Storage Area Network. All of this adds complexity and cost, as well as time delay, to the process of getting additional storage for the database.

Provisioning storage for a single database instance, let alone an entire enterprise, can be complex. Typically, the process runs something like this:

- Come up with a reasonable estimate of the amount of space you are likely to need
- Map out what you hope will be an optimal layout (placing data, log and archive files to avoid hot-spots)
- Create logical volumes
- Create file systems
- Install Oracle
- Create and load your database

And then the hard work begins: looking for hotspots that negatively affect performance, moving data files around to reduce contention, and dreading the day that you run out of space, and have to add more disks and shift all the files around again to rebalance across your updated storage configuration.

Fortunately, that scenario is about to change dramatically, with Oracle Database 10g's new Automatic Storage Management (ASM) capability.

ORACLE 10G PROVIDES A SOLUTION –AUTOMATIC STORAGE MANAGEMENT

In simple terms, Automatic Storage Management provides a vertically integrated files system and volume manager directly in the Oracle kernel, resulting in the type of ‘virtualization’ benefits (ease of manageability, higher availability) that storage area network systems provide, specialized for Oracle database files. ASM provides a quick way to lay out a database to accomplish optimal performance and utilization of resource. It enables a DBA to manage a pool of storage and lets the Oracle kernel manage the database files and placement of those files automatically. It removes the need for third party files systems and volume managers to manage Oracle files. ASM saves DBAs a great deal of time by eliminating the time consuming manual I/O performance and tuning tasks. It removes the need for the DBA to spend hours working with the System

Administrators to locate and remove hotspots in the underlying disk structure.

As ASM is an integrated file system and volume manager that has been purpose built to handle Oracle database files, it is not intended to be a general purpose files system nor general purpose clustered file system. In other words, it is not used for binaries, trace files and other non-Oracle files.

The process of laying out a database is much simpler with ASM than with conventional methods. Adding additional space is also greatly simplified with ASM. It is worth pointing out that ASM is not intended to be a tool for all layers of the provisioning process. ASM is used by the DBA to add disks to a disk group once it (block device or NAS filer) is recognized at the operating system level. The underlying steps of preparing storage to be recognized by the operating system still apply.

REMOVING COMPLEXITY

Automatic Storage Management eliminates the complexity associated with managing data and disks. It vastly simplifies the process of setting up a database, adding disks, and removing disks. Rather than managing hundreds, possibly thousands of files, DBAs using ASM create and administer a larger-grained object, the disk group, which identifies the set of disks that will be managed as a logical unit. The automation of the file naming and placement of the underlying database files saves the DBAs time and enforces best practice standards.

Adding a disk to an ASM disk group is as simple as preparing the disk for the underlying OS, and then issuing the 'add disk' command [`alter diskgroup <diskgroupname> add disk`]. The elegance of ASM really shows, though, when it comes to rebalancing the load after adding (or removing) a disk from a disk group. DBAs can use Enterprise Manager, DBCA or SQL*Plus to manage the disk groups. Whether you remove disks or add disks to your disk group, with a single command, all of your database files will be redistributed across the disks automatically, while the database is running. Because the algorithm ASM uses to spread data is not the rigid striping of traditional RAID, ASM does not need to re-stripe all of the data, but instead only needs to move an amount of data proportional to the amount of storage added or removed to evenly redistribute the files and maintain a balanced I/O load across the disks in a disk group.

In a configuration that is striped by a volume manager, the process of increasing the storage configuration often requires making additions of a minimum stripe set width or multiples of a RAID set. Then the process of spreading out the existing workload to take advantage of the new disks in the configuration may require exporting the database and reloading the entire contents into the newly re-striped storage configuration. The alternative is to carefully manage moving the database files around manually to shift the workload for optimal performance. This adds both cost and complexity to the process of improving performance. With ASM the rebalance process is very easy and happens without any intervention from the DBA or System Administrator.

ASM rebalancing distributes file data evenly across all the disks of the disk group. ASM automatically rebalances a disk group whenever disks are added or

dropped. Because ASM ensures that a file is evenly spread across all disks in a disk group when the file is allocated, rebalancing is not required except when the storage configuration changes. With the I/O balanced whenever files are allocated and whenever the storage configuration changes, the DBA never needs to search for hot spots in a disk group and manually move data to restore a balanced I/O load. As this is often a reoccurring event in a non-ASM environment, ASM saves significant time for the DBAs and System Administrators when looked at over the long term.

ASM eliminates much of the storage complexity associated with managing shared logical volumes that have often been used with RAC environments in the past. With shared logical volumes, the DBA has to coordinate with the System Administrator or Storage Administrator to create each logical volume that will be used as a database file. With ASM, the DBA can issue database commands to dynamically create database files in disk groups without involving any other administrators. This means that unlike shared logical volumes, disk group can be the destination for archive logs and RMAN backups. Furthermore, because ASM identifies the disks in a disk group by contents rather than by path, disks need not have identical OS paths on every node of a cluster. This adds flexibility while removing a significant amount of complexity from the process of provisioning storage for a RAC environment.

DISTRIBUTE I/O TO OPTIMIZE PERFORMANCE

Automatic Storage Manager spreads its files across all available storage for optimal performance. ASM enhances the concept of striping by providing additional flexibility for changes in the disk configuration. Spreading the workload across as many disks as possible results in less contention and high I/O throughput.

ASM divides files into 1MB extents and spreads the extents for each file evenly across all of the disks in a disk group. ASM does not use a mathematical function to track the placement of each extent; it uses pointers to record extent location. This allows ASM to move individual extents of a file when the disk group configuration changes without having to move all extents to adhere to a formula based on the number of disks.

For files, such as log files, that require low latency, ASM provides fine-grained (128k) striping to allow larger I/Os to be split and processed in parallel by multiple disks. At file creation time, a DBA can decide whether or not to use fine-grained striping. File type specific templates in the disk group determine the default behavior.

INTELLIGENT MIRRORING

Automatic Storage Management's native mirroring mechanism is an option that can be set in the process of creating a disk group. You can specify the redundancy protection for that group, and if you prefer to use an external mirroring system or you have no need for fault-tolerance, you can set the redundancy to *external*, in which case ASM will not mirror at all. Unless otherwise specified, redundancy defaults to *normal*, in which case ASM mirrors each 1 MB extent. For the highest level of fault tolerance, you can specify *high*

redundancy, which provides three-way mirroring.

If mirroring is selected for an ASM disk group, failure groups can provide an additional level of data. A *failure group* is a set of disks in a disk group sharing a common resource whose failure you wish to protect against. Once the DBA has configured a failure group, ASM will intelligently place the redundant copies of the data in separate failure groups to ensure that the loss of one failure group will not result in loss of data. The higher in the stack that one can enable redundancy protection the larger the failures they can survive. The use of failure groups enables larger loss of storage components to not result in any lost data.

HOW AUTOMATIC STORAGE MANAGEMENT WORKS

Automatic Storage Management has a special-purpose Oracle instance, called the ASM instance, dedicated to managing disk group activity. The ASM instance manages and communicates the map as to where each file extent resides. It also controls the process of rebalancing the placement of the extents when the storage allocation is changed. As an ASM instance uses only about 64-MB for its system global area, it requires a relatively small amount of system resource. In a RAC configuration, an ASM instance on each node in the cluster manages all disk groups for that node, in coordination with the other nodes in the cluster.

The ASM instance creates an extent map which has a pointer to each 1MB extent of the data file is located. When a database instance creates or opens a database file that is managed by ASM, the database instance messages the ASM instance and ASM returns an extent map for that file. From that point the database instance performs all I/O directly to the disks unless the location of that file is being changed. Three things might cause the extent map for a database instance to be updated: 1) Rebalancing the disk layout following an storage configuration change (adding or dropping a disk from a disk group), 2) Opening of a new database file and 3) extending an existing database file when a tablespace is enlarged.

An ASM instance cannot mount a database; it mounts disk groups and the database instance mounts the database. An ASM instance must be started before a database instance can access files in disk groups. Multiple and separate database instances can share disk groups for their files. On a single node, a single ASM instance typically manages all disk groups. In a Real Application Cluster environment, each node typically has one ASM instance that manages all disk groups for its node in a coordinated manner with the rest of the cluster.

All ASM management commands, such as creating disk groups or adding or dropping disks, must be directed to the ASM instance, and not to the database instance using the ASM Files. The status of disk groups and ASM attributes can be viewed through Enterprise Manager or through V\$views in the ASM instance.

USER TUNABLE/CONFIGURABLE PARAMETERS WITHIN ASM

The purpose of Automatic Storage Manager is to automate many manual storage management tasks. Like many other features in Oracle Database 10g

which help save DBAs time and effort, ASM is a self-managing capability within the Oracle kernel. That being said, there are a few parameters and configuration steps which need to be set when ASM is configured.

When creating a disk group the DBA needs to decide what level of redundancy ASM should provide. The choices here, as mentioned above, are external (no mirroring), normal (mirroring) and high (triple mirroring) redundancy. If mirroring is selected, the DBA can also assign disks within a disk group to be part of a failure group to provide additional level of protection.

Another parameter which the DBA can set is called the power limit. This is used to determine the priority that ASM should give to the task of rebalancing following a storage configuration change. This number can be set to any value between 1 and 11 with the default being 1. The higher the number the faster the rebalance will complete. The lower number the less I/O bandwidth is taken up by ASM to complete the rebalance, resulting in less impact on the database performance while the rebalance is taking place.

The ASM disk *discovery string* indicates where ASM should look for disks. When ASM mounts a disk group, it examines the headers of disks indicated by the discovery string to identify those which are part of the disk group. Although ASM has operating system specific defaults for the discovery string, other values may be defined to limit the number of disks examined in large systems or if a dynamic multi-pathing driver paths to block devices differ from those for default operating system block based devices.

INTEGRATED WITH ORACLE 10G COMPONENTS AND FEATURES

Oracle recommends using ASM with Oracle Database10g. Whether you are running on a single platform or in a RAC configuration, ASM works well with other new features in this release that involve managing the movement or placement of datafiles and space objects, including bigfile tablespaces, the disk cache recovery area, and Data Pump. As mentioned above, ASM is also integrated with SQL*Plus, Enterprise Manager, and the Database Configuration Assistant(DBCA). The ASM instance has been integrated with Recovery Manager (RMAN) and is required for backup and restore.

INTEGRATION WITH STORAGE VENDORS

As well as integrating with other Oracle functions, ASM is also integrating with many storage vendors. Several partners are working with Oracle to integrate their storage solutions with ASM, developing and extending APIs to further optimize the performance of an Oracle database running on their storage products. ASM has an interface that provides a means to prioritize I/O, enables caching hints, ensures data integrity before I/Os are committed to persistent media, and allows for more elegant vendor integrated disk discovery.

CONCLUSION

Oracle Database10g provides DBAs a simplified management interface for storage resources. Automatic Storage Management eliminates the need for manual I/O performance tuning. It virtualizes management of database storage to a set of disk groups and provides redundancy options to enable a high level

of protection. ASM facilitates non-intrusive storage allocation and provides automatic rebalancing. It spreads database files across all available storage to optimize performance and resource utilization. It is a capability that saves DBAs time by automating manual tasks and thereby increases their ability to manage larger databases and more of them with improved efficiency. In short, ASM is a feature that will simplify the job of the DBA.



White Paper Title
November 2003
Author: Paul Manning

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
www.oracle.com

Oracle Corporation provides the software
that powers the internet.

Oracle is a registered trademark of Oracle Corporation. Various
product and service names referenced herein may be trademarks
of Oracle Corporation. All other product and service names
mentioned may be trademarks of their respective owners.

Copyright © 2003 Oracle Corporation
All rights reserved.