
Oracle Database 10g's Automatic Storage Management Delivers Resilience, Simplicity and Lowers Costs

Date: August 2005
Author: Steve Duplessie and William P. Hurley
Title: Senior Analysts

Abstract: Oracle Database 10g's Automatic Storage Management (ASM) delivers previously unattainable levels of storage management efficiency through highly automated, application-centric rules and policies that reduce many hands-on tasks to the point of elimination. ASM enables continual access to database applications by delivering a near zero-downtime condition at the highest levels of performance. Oracle customers shared their experiences with ESG, enthusiastically affirming ASM's benefits.

Introduction

As the complexity of managing distributed business applications and their datasets increases, many organizations find themselves suffering from inconsistent application performance, unsteady availability and tedious storage and data management procedures – all at the price of higher operating costs. These performance, availability and operational challenges can cause IT organizations to miss SLA targets or force them to schedule frequent maintenance downtime periods all of which can hamper the effective execution of opportunities and obligations in a 24-hour business environment.

At the heart of most applications that drive and differentiate businesses of all sizes is a relational database management system (RDBMS). It is fair to say that relational databases have become the primary repository for business and mission-critical data. As such, mission-critical database systems must support both larger user populations and higher transaction volumes all while meeting service level agreements requiring the highest levels of responsiveness and availability.

Businesses' ever increasing reliance on more – and more sophisticated – database systems has contributed to a dramatic growth in the capacity and complexity of data storage sub-systems. Data growth and storage system complexity are further accelerated by obligations to comply with new government regulations and mandates to maintain auditable records for extended periods – even decades. Even though breakthroughs in storage hardware technologies have reduced the purchase price for arrays of ever-larger capacities, management tools for these disparate heterogeneous systems has lagged. Database, storage and system administrators are frequently left with software tools that can hamper the effective and efficient management of the storage systems that house this growing pool of vital data.

To overcome the challenges described above, Oracle has developed Automatic Storage Management (ASM). ASM is a storage and data management capability that is built into the kernel of Oracle Database 10g. ASM provides database administrators (DBAs) system administrators and storage administrators with new capabilities and automated features that can both lower the cost of storage ownership and increase the quality of service delivered by Oracle

database systems. ASM is an integral part of Oracle Database 10g and is provided free-of-charge with all Oracle Database editions.

This paper will review the technical attributes of ASM and will discuss the benefits of employing this capability. Our discussion of benefits will be based on interviews with IT Professionals such as with database, system and storage administrators who are using ASM today and will be complemented by our own research.

Challenges of Data and Storage Management

Until recently, product vendors as well as end-users viewed data management and storage management as distinct though overlapping functions. Even though this view still predominates, it is an artifact of the computing and application models that succeeded in replacing fully integrated mainframe systems. The various components of the post-mainframe model required specially designed management tools to operate hardware systems, configure or provision logical resources, manage discrete databases and to manipulate data (e.g. replicate, back-up or prune) without corrupting its integrity or accessibility. This complement of requirements led to the need for management tools that were either system specific or designed to operate on a homogenous “stack”.

The proliferation of post-mainframe “open systems” technologies spurred vendors to deliver an assortment of new management tools. These tools generally addressed the needs of the database, system and storage administrators but also reinforced the further division of skilled IT Professional who were responsible for an increasingly complex, yet stove-piped infrastructure. Under these conditions, when a DBA needed more storage capacity they would issue a request that both system and storage administrators would work to fulfill. DBA's would then be presented with a slice of storage capacity (LUN) that might be integrated into an existing logical data volume. Frequently, the entire team would be required to participate since adding new capacity required application downtime, data migration, volume reconfiguration and updates to security and access mechanisms.

The advent of storage networking changed the paradigm for designing and operating data storage infrastructure that was supporting of two-tiered, “open systems” architectures. Prior to storage networking, server systems could be clustered to provide higher levels of system resilience and performance. Yet, these models too required data and storage management tools that differed from standard two-tier models. Vendors soon delivered solutions that allowed end-users to wed clustered server systems to networked storage, primarily across a fibre channel storage area network (SAN). The arrival of SAN and other storage networking technologies created new opportunity to mix and match various hardware system resources by exploiting server clustering, SAN and other network infrastructure. By mixing and matching resources, IT professionals could deliver higher levels of resource utilization by storing data of high-value on high-function systems and data of lower-value(s) on lower-function arrays. Even though administrators could network storage systems of various types and sophistication, most data and storage management tools were not adequate to address hardware heterogeneity and data structure differences.

Despite the real improvements delivered by the solutions described above, database, system and storage administrators frequently work in isolation. Cross-functional actions or integrated system support efforts were only undertaken after long hours of planning and interdepartmental coordination. The occurrence of a real-time problem presented these IT professionals with extraordinary logistical challenges on top of the immediate technical issue. Even with these extra efforts and advanced storage systems, IT professionals still struggle to realize higher degrees of freedom in mixing and matching resources, fully and reliably utilizing automated functions and reducing the number hours dedicated to planning, integrating and rolling-out application enhancements and expanded data access.

Fortunately, some vendors are now providing solutions that address the full range of challenges described here. As this paper will demonstrate, Oracle Database 10g with its native Automatic Storage Management capabilities enables IT professionals across the support and operations value chain to achieve many of the operational outcomes.

Introducing Automatic Storage Management

Automatic Storage Management is a new feature in Oracle Database 10g that provides data management features within the database that simplifies the management of database systems and storage environments. ASM is the combination of a clustered file system and logical volume manager that is tightly integrated and embedded within the Oracle database 10g. ASM was specifically designed to operate with Oracle database files. ASM is a standard Oracle

Database 10g feature that is included at no additional cost with the database system. ASM is designed to optimize the operation of Oracle databases and simplify storage management by virtualizing storage hardware, automating tasks and optimizing performance thereby reducing operational overhead, administrator intervention – and human error – and overall management costs.

ASM is application aware, meaning that it understands the content of the files it manages. This awareness enables ASM to rebalance and redistribute data in ways that a separate logical volume manager cannot. DBAs can achieve performance levels that are similar to using raw devices because ASM essentially eliminates the need for a third-party volume manager and supersedes the file system of host operating systems.

At the heart of ASM is the “ASM Disk Group.” A disk group is an abstraction of physical storage media that virtualizes the access to and management of these hard-disk systems and the files stored on them. ASM disk groups are defined by the DBA who sets the definition of the physical disks. The Oracle database kernel manages the database files stored within that disk group pool. Furthermore, storage capacity connected to that disk group can be added or removed. All interaction between the Oracle databases in the storage media is implemented through these disk groups, eliminating the need to use overlaid third-party volume managers or file systems to manage Oracle database files.

ASM is Oracle’s collection of intelligence and best practices learned over many years. The ASM disk group virtualization interface allows DBAs to easily and dynamically add or remove disks for a particular disk group without interrupting the database operation. ASM will automatically begin to use a new disk as soon as it is added to the disk group, and will automatically rebalance data between old and new disks to continually optimize I/O performance. ASM also eliminates disk fragmentation and hotspots thereby reducing management overhead.

ASM provides up to 3-way mirroring for higher availability. This mirroring function is especially important where low cost storage solutions are deployed. The lower the storage array functionality, the higher the value of ASM in terms of providing striping and mirroring. The Oracle DBA simply defines a failure groups within a disk group and ASM will automatically mirror the data across two or more failure groups, which can help to avoid downtime in the event of component failures, such as a disk controller or even entire disk array. This powerful redundancy feature increases availability without any additional management overhead. Since failure groups are automatically managed by ASM fail over or fail back activities are fully managed without the intervention of an operator by ASM when the failed component is replaced.

ASM can mirror and stripe data across multiple disks spindles and can dynamically redistribute data among those disks to spread the I/O load; all of this can be managed from within the database using Oracle Enterprise Manager and Grid Control.

Customer Results

Customers that have deployed ASM are extremely enthusiastic about its capabilities. These customers have realized improvements across a number of functional and operational areas.

Mainstream Technologies

One customer we interviewed indicated that his organization experienced “a complete paradigm shift” in how it used Oracle and ran IT operations “after implementing ASM.” Mainstream Technologies is a managed service provider that operated and maintains complete Oracle infrastructure environments for its customers. Mainstream customers are almost evenly split in number between those that outsource the hosting of their Oracle infrastructure behind the walls of Mainstream’s data center and those that keep it in-house.

Prior to adopting Oracle Database 10g and implementing ASM, Mainstream had many challenges keeping its customers’ Oracle environments current and effectively tuned in line with SLAs. Since Mainstream’s customers tend to be geographically dispersed, careful planning of scheduled maintenance and tasks was necessary to the distances between on-site locations. If issues arose at customer sites, it was both costly and time-consuming to dispatch service professionals. In many instances applying updates, modifying procedures, changing queries, adding capacity or other maintenance tasks required application or system downtime. Those challenges directly impacted Mainstream’s SLAs, proving costly to both Mainstream and its customers.

Mainstream adopted Oracle Database 10g because of ASM in order to address these issues with its customers. Oracle Database 10g management framework geographically dispersed sites to be managed from a central location. As ASM is integrated into the central management framework, Mainstream could now manage its data and storage SLAs across all customers from its own headquarters, essentially eliminating the need to dispatch service professionals. One of the most common tasks Mainstream was able to simplify through automation and centralized control was the rebalancing of database hot spots. Previously, this activity was undertaken manually by first running a series of test scripts to determine usage patterns and data layouts, then – after quiescing the database – the on-site team would manually copy and relocate files on the storage sub-systems in order to improve performance and meet SLA targets. Often, this procedure could take an entire weekend. Now, Mainstream uses to ASM proactively scan for hot blocks or spots and automatically rebalance the data layout in real-time. This function alone allowed Mainstream to eliminate 60% of the human resources required rebalancing data sets and meeting its SLAs.

Mainstream also utilizes ASM to quickly and efficiently add storage capacity in line with SLAs that define the “class-of-storage” required for a given application. Mainstream uses ASM Disk Group functions to define high, medium and low value storage to be used or assigned to the appropriate storage class to a specific application. Since applications have different availability, scalability, reliability and performance requirements, ASM allowed Mainstream to these requirements with the most appropriate class of storage. Mainstream not only benefited from exact storage type to requirement mapping in service of SLAs, it also lower the overall cost of acquiring storage since Mainstream IT Professionals could now reliably intermix expensive, highly functional arrays with less expensive alternatives with no penalty for securely providing data mirroring and protection.

Mainstream has implemented a great deal of automation through ASM, which they believe helped them become a trusted partner to their customers. By implementing these automated features, Mainstream could offer more services with far greater economies of scale allowing it to achieve a much larger market footprint with the same staff levels.

Talk America

Another customer we spoke with, Talk America, a telecommunications provider with retail, commercial and private-label customers. The company uses Oracle Database 10g for the majority of its applications such as billing, customer support, customer self-service, data warehousing, CRM and other business processes. Before migrating to Oracle Database 10g with ASM, the company had 10 Informix systems each hosting its own database and with its own storage resources. After migrating to 10g with ASM the company consolidated database instances and move to a SAN-based shared storage infrastructure. Using ASM to consolidate and share resources delivered immediate and significant software license saving to Talk America.

By moving to a shared storage environment, Talk America was able to effectively and reliably allocate storage by class to each of its 11 main business applications. As with Mainstream, Talk America is now able to intermix different storage arrays then logical carve and provision classes of storage matched to the requirements of a workload or application SLA all through ASM and with the “disk group” function. ASM also allows Talk America to automatically reassign or re-provision storage on a rules and policy basis. In instances where a Business Intelligence application needs more capacity for a data warehouse or mart, ASM can identify unused capacity within the disk group pool and automatically reassign that fallow space to a meet a real-time need – all with no operator intervention.

The simplicity ASM enables when expanding an existing data volume and adding capacity delivered immediate operational benefits to Talk America. Previously, undertaking such a task required input from storage and system administrators and, most likely, application downtime. The old procedure was driven by a 2 Gigabyte LUN capacity restriction that led to data volumes that consisted of a hundred or more LUNs being rolled up under one 3rd party logical volume manager. With ASM, Talk America IT Professionals simply draw the necessary capacity from a spare pool and add it to existing disk groups – all in real-time, all with no downtime of application quiescence. On top of the benefits of streamlining operations and multi-functional procedures, Talk America significantly reduced the cost and complexity of 3rd party software licenses, maintenance and staff training.

Talk America's business is growing rapidly and as such its storage needs are dynamic. ASM gives the company the ability to react quickly to new business needs and to fully automate formerly labor-intensive provisioning, allocation and optimization tasks. Even though these tasks may be automated through extensive scripting or other 3rd party tools, it is ASM's ability to perform these and other capabilities proactively and without application downtime that allowed Talk

America to simplify its IT operations and reduce cost. With ASM in place, Talk America can now perform almost any business process in real-time.

A clear-cut and beneficial example is billing reporting. Formerly, billing reporting was done on a batch basis on a daily or weekly as needed. Data from the various independent systems was extracted and loaded into the Talk America billing data-mart for processing. Frequently, spare storage capacity was needed temporarily in support of the billing batch job. Allocating this temporary capacity was done after hours in a labor-intensive fashion. This interval-based billing process limited Talk America's ability to offer a more comprehensive suite of services, including lucrative pre-pay and self-service offerings. Now that ASM is in place, Talk America draws upon billing information in real-time thereby delivering superior customer support and self-support as well as a complement of new services that is stimulating the growth we mentioned at the outset.

The Bottom Line

Many storage companies sell separate software products to accomplish auto-provisioning and storage "tiering," but these products are not tightly integrated with the application. Using data and storage management tools that are tightly integrated with and aware of an application's behavior makes sense because it is through applications that data is created or consumed. Our research, bolstered by anecdotes from the field, indicates that this level of integration and awareness leads to high performance, more reliable automation and simplified operations – as achieved with ASM.

Additionally, ASM's real-time capabilities differentiate the product, Oracle and customers using the product. Real-time capabilities such as automatic rebalancing and database optimization, capacity provisioning and multi-way mirroring allows end-users to more fully utilize storage resources while providing thorough data protection measures across a variety of storage hardware types. The efficiencies facilitated by ASM storage virtualization enables application-driven tiering of by storage by class, allowing IT Professionals to mix and match low- and high-function storage arrays within one common framework. Lastly, ASM's real-time capabilities have clearly enabled end-users to reduce operator intervention and overhead – by as much as 60% - but to also fully realize the upside potential of today's distributed application architectures.

The Oracle database has arguably always been the most important storage management product ever created, and with ASM, it extends its lead even further. The most valuable data in the world lives and breathes inside Oracle instances in the largest companies on the planet – and ASM will help improve the performance and manageability of that data.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. and is intended only for use by Subscribers or by persons who have purchased it directly from ESG. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.