

DATABASE SERVER MANAGEABILITY STRATEGIES AND DIRECTIONS

Richard Sarwal & Sushil Kumar, Oracle Corporation

INTRODUCTION

The Oracle database has always been known as the industry leader in terms of performance and reliability, while at the same time providing extremely rich functionality. Most successful businesses today use the Oracle database to power their mission critical applications. However, as the business environment becomes increasingly competitive, the challenge before today's enterprise is to manage their information technology (IT) infrastructure at the lowest possible cost without compromising service level objectives. This presents an interesting dilemma since as systems provide more and more capabilities system complexity can increase, and hence these systems can become expensive to manage and sustain. Recognizing these demands, Oracle has done significant work to improve the manageability and ease of use of its products. This paper discusses Oracle's strategy, directions, and technology to make the Oracle Internet platform inherently simple and self-managing so that the total system deployment cost is reduced.

KEY MANAGEABILITY CHALLENGES

In today's round-the-clock Internet economy, the importance of an efficient and reliable IT infrastructure for the success of an enterprise hardly needs any explanation. As businesses increasingly rely on this infrastructure to service customers and partners, and disseminate information among employees, computer systems are no longer peripheral to the management of core business. Even a temporary outage of a critical application can jeopardize the viability of the entire business. Revenue and customers may be lost, penalties could be owed and, the resulting unfavorable publicity can have a disastrous affect on customer confidence and corporate stock valuation. The effective management of the enterprise IT infrastructure is therefore key to the success of modern enterprises.

Recognizing this need, companies today are investing significant amounts of financial and administrative resources to ensure the best possible management of their IT infrastructure. However as digital gadgets become increasingly pervasive in our daily lives, organizations are forced to deal with more data than ever, with customers and employees needing around the clock access to this data, from more places, via more type of devices. As IT vendors deliver increasingly sophisticated solutions to meet such exacting demands, the task of systems management has never been more complex. Hiring highly skilled administrative staff to manage such complicated environments is an expensive proposition. This coupled with frequent shortage of experienced administrative personnel often results in spiraling management costs.

While adjusting to new technologies and business practices presents numerous technical challenges, businesses today are also faced with increased competition as the Internet has collapsed the world into a small global village. An online bookseller in Seattle faces as much competition from other booksellers in the US as it does from one in China. In order to maintain business profitability amidst ever-growing competitive pressure, corporations must minimize their operating expenses with system administration costs being no exception. This creates an interesting corporate challenge. Enterprises must manage their systems much more effectively than ever -- to ensure the highest performance, scalability and availability – but at a cost significantly lower than before.

ORACLE'S MANAGEMENT METHODOLOGY

Rising to these challenges, software and hardware vendors all around the globe, including Oracle, are trying to create products that require as little manual administration as possible. However, what sets Oracle's approach apart is the breadth and depth of its scope. Unlike other vendors who are focusing their attention on individual products, Oracle is looking at this problem from an overall business perspective. A typical corporate information management system

today comprises a set of diverse software and hardware products integrated together with pieces of data spread across a number of data centers all over the world. Solving the manageability problems of the individual components of the enterprise IT stack may assuage the pain a little bit but it certainly does not solve the bigger problem since customers still have to deal with the challenges of putting the pieces together. Oracle's management methodology, therefore, encompasses a broad based initiative which not only seeks to minimize administrative costs by making Oracle products completely self-managing but also tries to holistically address the technical challenges faced by IT managers and administrators in creating and managing the enterprise eBusiness solution.

SIMPLIFIED AND INTEGRATED PLATFORM

Studies show that companies spend approximately 80% of their IT budget to build and maintain the standard computing infrastructure, and only 20% for creating competitive differentiation. Why does it take 80% of the IT budget to provide standard business services? The reasons are not hard to find. Developing and deploying web applications usually involves integration of several disparate components, each of which needs to be administered differently. Also, the underlying data is stored in many different places such as databases, multi-dimensional OLAP repositories, file systems etc., further increasing the administrative complexities. Integrating these diverse products, that were not designed to work together, into a consolidated set of business applications is difficult to do, difficult to support, and therefore, highly expensive.

By integrating all its infrastructure products under the Oracle9i umbrella, Oracle has provided an extremely cost effective means of building and deploying dynamic web applications. Oracle9iAS is a comprehensive, integrated application server that provides all of the infrastructure and functionality needed to run every successful eBusiness. It is a deployment platform that supports many different development approaches and provides a market leading implementation of the latest and emerging industry standard technologies including Servlets, JSPs, EJBs and XML. All components of the Oracle9iAS technology stack are tightly integrated and have been extensively tested for seamless interoperability thereby eliminating the need of administering them individually. Administrators, therefore only need to install and manage a single product to support all their enterprise Internet application needs – be it a web site, a portal or a web based business application.

While the common application infrastructure provided by Oracle9iAS significantly reduces integration costs, the Oracle9i database server enables businesses to consolidate all their data into a single data store. The Oracle9i database stores and manages not just characters, numbers and dates, but also location information, email, documents and files, and media such as image, audio and video. In addition, the Oracle9i database includes the ability to integrate an organization's existing third-party information sources through transparent gateway technology. By providing an integrated and complete solution for all business information, the Oracle9i database minimizes the cost of storing and using all the different types of business information modern organizations have come to depend on.

Both Oracle9i Application and Database server fully leverage each other's capabilities. The ability to use common programming languages (SQL, PL/SQL, Java, XML), APIs (OCI, JDBC, SQLJ, XDK) and development tools allows all database developers to become productive web developers, and vice-versa, with absolutely no additional training. The complete Oracle Internet platform can be administered from a common console thanks to enhanced Oracle Enterprise Manager product suite. From client to application server, database, host and network, Oracle Enterprise Manager has automated and centralized the management of a completely integrated business infrastructure, resulting in higher administrator productivity, enhanced quality of service and greatly reduced operating costs.

CERTIFIED CONFIGURATIONS

Building an enterprise computing infrastructure also involves setting up the operating environment by assembling a multitude of components, such as the server hardware, operating system, storage subsystems, utilities etc. Traditionally, the task of integrating these components has been left to customers. Administrators often find themselves struggling with issues like: *which version of cluster manager software do I need to use with a given version of the OS and database software?*, OR *do I have all the required patches for different software components in order to ensure uninterrupted operations?* These are some of the extremely critical decisions that an administrator needs to make, since failure to correctly configure the system may result in compromised security, availability and performance. Yet in the absence of any clear guidelines, the administrator may be left with no option but to learn from his/her mistakes.

Oracle, therefore, has taken the lead in collaborating with other industry leaders to define new standards for quality and interoperability. Working with its partners, Oracle is developing fully integrated solutions to meet the widest spectrum of business requirements. As a part of the VOS Initiative (<http://www.vosinitiative.com>), Veritas, Oracle and Sun are working together to create unified solutions that are pre-tested and validated for interoperability and reliability. A jointly staffed escalation center has been created for effective interoperability support. Customers using these unified solutions may contact any of the three companies when experiencing a problem. If the problem is identified as an interoperability issue, it will be automatically escalated to the joint escalation center. This along with joint consulting services and co-developed training solutions help customers deploy complete IT solution faster, with great reliability and, with quicker resolution to complex problems.

Similarly, under the ECOstructure initiative (<http://www.eECOstructure.com>), Oracle, EMC and Cisco are combining their solutions and services to deliver an integrated, tested and validated eBusiness infrastructure suite. The three companies have committed joint engineering and support resources to drive the development of combined technology blueprints. These blueprints contain guidelines for designing, implementing, and managing a highly scalable and reliable infrastructure based on Oracle, EMC and Cisco technologies. A joint development center tests and validates the blueprints as well as measures their response to scenarios that could limit availability, data protection, scalability and security.

In addition to the programs highlighted above, Oracle continues to work very closely with a number of other partners to seamlessly integrate the Oracle database server with a wide variety of supporting tools and technologies.

DATABASE CONSOLIDATION

The continued explosion in the volume of data being generated by today's businesses has led to proliferation of small departmental/regional data centers all across the world. Each of these data centers typically manages one or more local databases. The resulting data fragmentation could adversely impact the effectiveness of the business management process as managers may not be able to access the consolidated information when needed. In addition, it creates significant inefficiencies in IT operations since a number of administrative tasks must be performed multiple times, once for each database. Consequently, as the number of databases increases, so does the management complexity.

The latest technological advances provide every incentive for consolidating these smaller databases into either a single consolidated database or a few large databases depending on organizational constraints and business requirements. Using high-end enterprise server hardware available today augmented by the rich functionality provided by the Oracle database for scalability, security and resource management, it is now possible to host even the largest of databases on a single node without compromising business objectives. For example using features like Virtual Private Database and Oracle Label Security, it is possible to implement a fine-grain access control mechanism that prevents unauthorized access even for data stored within a table. Different organizations within a company can therefore keep their data in a common database but can only access the data that pertains to them. Similarly resource management products, such as OS resource management tools as well as Oracle Database Resource Manager, allow administrators to distribute shared resources among different users and applications according to business priorities. This approach to resource management is vastly superior to static resource partitioning since any unused resources can easily be made available to applications that need them. This allows for greater flexibility in resource deployment thereby minimizing the wastage caused by the inability to share the "compartmentalized" resources.

While the technology provides the impetus for consolidation, the biggest incentive comes from enhanced operational efficiency and reduced administrative costs. The consolidated databases can be managed by a significantly smaller administrative staff with far greater effectiveness, which not only drives the operational cost lower but also results in improved service. Oracle has employed this strategy to optimize its IT operations with phenomenal success. A worldwide network of 97 email servers and 120 databases has been consolidated into 2 servers running 4 databases. Leveraging the enhanced efficiency gained through consolidation, the centralized US data center is now being managed by 100 fewer administrative personnel while supporting 10,000 more users. As a result, Oracle's global IT operating budget has been reduced from \$600 million to \$400 millions realizing a saving of \$200 millions annually. At the same time, quality of service has significantly improved thanks to enhanced administrative productivity. Oracle

Global IT is today providing more services to Oracle employees worldwide with better availability and reliability-- but at a significantly reduced cost. *To learn more about Oracle's success story with database consolidation, please refer to Oracle white paper "Keep It Simple: How Oracle Consolidated Its Global Infrastructure into a Centralized e-Business Architecture" at http://www.oracle.com/collateral/global_it_white_paper.pdf.*

SELF-MANAGING SERVER

The Oracle database has been in existence for over twenty years with every new release providing higher levels of functionality. However, while adding new capabilities to the database, Oracle has simultaneously taken steps to eliminate complexities within the server in order to simplify its administration. The introduction of version based read consistency in Oracle 6 established an incredibly transparent mechanism of managing the undo (Rollback) data. As opposed to storing undo data in the recovery log file, as done by most of the competing products such as DB2 and SQL Server, Oracle stores this data right inside the database in structures called undo segments. This ground breaking technology allows the Oracle database to service queries without requiring any read locks as well as to wrap around the log file without waiting for transactions to commit. Consequently, Oracle DBAs need not even think about some of the most time consuming tasks that their DB2 and SQL Server counterparts have to perform on a daily basis i.e. resolving deadlocks and monitoring uncommitted transactions.

Version 7 of the database server continued to enhance manageability by introducing a number of features to simplify day-to-day administration. Among other things, this release introduced several new wait statistics to facilitate troubleshooting and, simplified the management of the shared memory by consolidating various library cache and row cache sizing parameters into a single global control called the SHARED_POOL_SIZE. A new graphical tool, Oracle Enterprise Manager, provided administrators with an extremely simple and easy to use interface for performing common database management tasks. To meet the growing challenges of managing large volumes of data, Oracle8 introduced the concept of partitioned tables. This feature allowed administrators to sub-divide a large table into a number of smaller partitions which could be managed independent of each other. Routine maintenance operations such as data archival and data reorganizations could be performed on individual partitions minimizing their impact on overall system availability and performance. This release also introduced a new utility, called the Recovery Manager, to streamline the management of backup and recovery operations.

Locally Managed Tablespaces in Oracle8i eliminated the need for periodically reorganizing tablespaces to reclaim fragmented space, thereby eliminating a major task from a DBA's to-do list. Oracle8i also unveiled a new resource management functionality to enable implementation of resource allocation policies within the database server. Using Oracle8i Database Resource Manager, administrators are able to distribute CPU resources among database users and applications according to business priorities. By partitioning system resources among tasks of varying importance, the Database Resource Manager made it possible for administrators to guarantee measured database services to enterprise users and applications.

The latest release of the Oracle Server, Oracle9i Release 1, has been a significant release from the manageability perspective. With enhanced manageability being one of the key focus areas for this release, Oracle9i has automated a number of routine DBA tasks, reduced complexity of administration and made the server more self-tuning. Thanks to Automatic Undo Management, administrators no longer need to manage the rollback segments manually. A new advisory allows administrators to size the buffer cache optimally by predicting the number of physical reads for different sizes of the cache. With Resumable Space Allocation, it is now possible to intervene and correct errors in the middle of an operation execution and thus prevent space related failures. Features such as Automatic Segment Space Management and Automatic SQL execution memory management have automated some of the most difficult and time-consuming performance tuning operations. The advances made by Oracle9i in the manageability sphere have been widely applauded by both customers as well as industry observers. 170 Systems (<http://www.170systems.com>), a web-based solution provider for business procurement enhancements and content management, finds Oracle9i to be an excellent platform to build, deploy and support their solutions. According to Karl Buttner, President, 170 Systems, "Enhanced administration and database management capabilities allow our customers to reduce their administration and operational costs, while improving availability and performance of their systems". Mark Shainman, a META Group senior research analyst, concurs – "While Oracle has always included significant functionality in its DBMS product, it generally requires various "tuning" to achieve optimal performance. (Oracle)9i automates many secondary processes, enabling easier management and DBA user friendliness". To learn more about how customers and partners are benefiting from the integrated management

solution provided by Oracle9i, please log on to http://oracle.com/ip/deploy/database/oracle9i/index.html?ma_home.html.

WHAT'S NEXT?

ORACLE'S MANAGEABILITY VISION

While Oracle9i has come a long way in simplifying manageability challenges, businesses still find database management to be a complex task. There are too many concepts to be learnt, too many commands to deal with, and way too many statistics to analyze. Oracle has, therefore, made manageability a strategic development area. A core group of senior developers has been put together to spearhead the Oracle manageability initiative – a highly ambitious project with the objective of building a completely self-managing database server. The charter of the group includes carefully investigating all the administrative tasks that a DBA needs to perform currently and exploring the possibility of automating them. The vision of the group goes beyond just introducing a set of new manageability features. It seeks to take a comprehensive look at the server architecture in order to identify areas where:

- Existing complexities could be eliminated
- New components can be built to enable administrative automation
- The administrative interface could be significantly simplified

In addition, a set of guidelines is being developed to ensure that all future development work includes manageability as a key objective. The end goal of this project is a product which can be installed and configured by answering a few simple business related questions such as performance and availability requirements, the nature of the application accessing the database, the maximum amount of system resources that the database is allowed to use, etc. Once the configuration is completed, the database would monitor the workload and automatically adjust its configuration to meet the stated business requirements. For problems related to resource scarcity, or other such issues that can't be resolved without human intervention, the server will automatically send a notification to the administrator along with the suggested remedy. It will also have the intelligence and knowledge to answer common capacity planning questions such as how many users can my system support with existing resources without compromising performance goals, how will the system be impacted if I were to add or remove some hardware resources, or how much more resources do I need in order to support a more aggressive set of business goals, etc. To sum it up, Oracle is seeking to create a product that understands business requirements, manages itself automatically and, satisfies the even the most stringent performance, scalability and availability goals out-of-the-box. At the same time, it should provide enough flexibility to let an expert administrator take over if desired. The need to do so, however, should be extremely rare and completely optional.

This, obviously, is a very challenging set of goals involving finding solutions to some of the long unresolved computer science problems. Since the Oracle database server runs under every popular operating environment, it also requires very close collaboration with the OS and hardware vendors. As such, no definite commitments can be made about whether, when and how these solutions will be implemented. However, continuing its pioneering work in the area of database software development, Oracle has taken the initiative to develop a complete manageability solution putting aside all cynicism about the feasibility of such an exercise. The following sections provide more details about the strategies and directions adopted by Oracle for making the Oracle database server completely self-managing.

INSTALLATION & CONFIGURATION

Problems encountered during installation can be extremely frustrating. It is no exaggeration to say that the usability of a product is often gauged by how simple or complex its installation is. Oracle is, therefore, investing significant amount of development resources in simplifying the installation, configuration and upgrade tasks. The Oracle Universal Installer (OUI) is being enhanced to automate all pre and post installation tasks. In the future, OUI will automatically check the system prior to beginning the installation to ensure that the OS has been configured properly, the correct patches have been applied, and enough resources are available to guarantee the success of the installation process. If any problems are detected during the pre-install check, the installer will help the administrator fix the problem by recommending corrective steps. For example in an UNIX environment, if the installer detects that the

kernel shared memory parameters have not been properly configured, it may recommend that administrators execute a supplied script to correctly set up the OS kernel parameters. Also, the installation process will be completely self-contained and will automatically set up the required infrastructure for routine monitoring and administration. Tools such as Enterprise Manager and Recovery Manager (RMAN) will be automatically configured and, all Oracle server components such as the database, listener, management framework etc. will be configured for automated startup and shutdown. All the development work currently underway to make the OUI more robust, intelligent, and fault tolerant will make it possible for anyone with a basic familiarity with computers to install and get started with the Oracle database server.

Oracle may also be able to fully automate the process of managing software patches in the future. It may be possible to automatically identify the require patches and download them to a location suggested by the administrator. The possibility of automatically applying the downloaded patches during an administrator specified time interval is also being explored.

The database creation process is being simplified by making the Database Creation Assistant (DBCA) smarter. While the equivalent command line interface will continue to be supported, all future manageability enhancements will be restricted to this tool. In the future, DBCA will allow administrator to create all possible configurations of the database be it a stand-alone database, a Real Application Cluster (RAC) database, or a standby database. During the database creation process, the DBCA will guide administrators in setting up a backup and monitoring plan. The newly created database will also be automatically registered with centralized Enterprise Manager, Recovery Manager, and LDAP repositories. A database created using the DBCA will, therefore, be fully setup and ready to use in all respect. Similarly, Oracle Database Migration Assistant (ODMA) is being enhanced to make the database upgrade/migration process extremely simple. Using the next generation ODMA, administrators will be able to migrate/upgrade any database configuration, including RAC and standby, just by answering a few simple questions. The tool will automatically check the system if adequate resources are available, ensure adherence to the best practices – such as backing up the database before beginning the upgrade/migration process, replacing the obsolete and deprecate initialization parameters etc – and, verify the successful completion of the operation. In addition, it will also help administrators in activating the new features available in the release being upgraded to.

The Oracle database server provides a number of initialization parameters to optimize its operation in diverse environments. Only a few of these parameters need to be explicitly set as the default values for the rest of them are adequate in vast majority of cases. However, competitors have been trying to create confusion by claiming that the large number of Oracle initialization parameters presents a serious manageability challenge for customers. Oracle is, therefore, contemplating categorizing the parameters into basic and advanced categories. Administrators will be able to restrict their day-to-day interaction with the basic parameters whose number is unlikely to exceed 25. The advanced parameters will allow expert DBAs to adapt the behavior of the Oracle database server to meet unique requirements without overwhelming those who have no such requirements. The future releases of the Oracle database server will, therefore, provide the right blend of simplicity and flexibility – simple enough to be used as an embedded database yet flexible enough to meet even the most challenging requirements.

ONGOING ADMINISTRATIVE TASKS

STORAGE & SPACE MANAGEMENT

Database space management has always been an important part of any database administrator's job. Administrators spend a significant amount of time planning and monitoring the space utilization in order to ensure uninterrupted database operations. Current research at Oracle is therefore directed towards automating the storage and space management activities within the server making them completely transparent to administrators. The Oracle Managed File feature in Oracle9i has already automated the management of database files. While using this feature, database administrators no longer need to directly manage the files of an Oracle database since the server automatically creates and deletes them as needed. The new technologies being developed will allow the Oracle database server to completely automate the task of storage management as well. Administrators will be just required to specify a set of disks that an instance is allowed to use. Oracle will then automatically determine the most optimal storage layout and rearrange the data, as and when required, to ensure best performance.

The server will also monitor space utilization and notify administrators if the database is likely to run out of free

space in the near future. It will detect data fragmentation and automatically rebuild the objects online without impacting the system performance. Other similar tasks such as coalescing sparse indexes, rebuilding fragmented and stale indexes and, removing chained rows may also be automatically performed by the server. Allowing the datafiles to grow and shrink automatically will lead to more effective space utilization, thereby ensuring that the administrator intervention is only required when the system is truly out of disk space.

As explained earlier, Oracle has simplified the space management tasks considerably over the years by introducing features such as Locally Managed Tablespaces and Automatic Segment Space Management. These features eliminate the need of performing any manual tuning, provide better out of the box performance and, improve the space utilization. With the availability of locally managed SYSTEM tablespace in Oracle9i Release 2, the stage is all set for the complete desupport of the deprecated features such as dictionary managed tablespaces and manual segment space management in the near future.

Continuing the same trend, Oracle is working toward complete automation of the undo and temp space tuning. The introduction of Automatic Undo Management in Oracle9i liberated administrators from adjusting the attributes of rollback segments to avoid undo block and consistent read contention. This feature also gave administrators control over the retention of undo data by allowing them to specify the undo retention period in terms of wall clock time. Future enhancements to the Automatic Undo Management feature may allow the database to automatically adjust the retention time based on the nature and length of queries accessing the database. The granularity of undo retention may also be changed from the database to individual object level allowing the server to retain the undo data for frequently accessed tables longer than that of inactive objects. Temporary tablespace quotas may be introduced to provide a temporary space reservation mechanism for critical operations eliminating the need to create multiple tablespaces. The server may also be able to prevent failures in case of temporary tablespace shortage by delaying, rescheduling or preempting execution of certain operations. Enhancement such these will simplify a DBA's life considerably by handling the space pressure situations in a graceful manner. They will also ensure that the only day-to-day space management task left on a DBA's task list will be to monitor the system disk space availability!

BACKUP & RECOVERY

Oracle's current backup & recovery infrastructure is already considered the most robust and powerful in the industry. Oracle Recovery Manager (RMAN) is a powerful tool that simplifies, automates, and improves the performance of backup and recovery operations. Though Oracle continues to support user managed backup & recovery, all future manageability enhancements will be restricted to RMAN. Oracle strongly recommends that customers use this tool since it provides the most efficient and error-free method of performing backup and recovery operations.

Oracle9i Release 1 included many enhancements to make RMAN easier to use and more self-managing. Oracle9i RMAN automatically manages backup retention as per an administrator-defined policy. It can also resume a failed backup or restore operation thereby eliminating unnecessary duplication of work. These, and other such improvements in Oracle9i, have made RMAN extremely easy to use and highly fault tolerant. Enhancements currently underway may enable RMAN to completely automate the database recovery process. Using a single command, administrators may be able to direct RMAN to automatically detect the nature of the failure and repair the database using the most efficient recovery method. Also, administrators may no longer need to manage the location of different types of recovery files, such as archive redo logs and on-disk backups, manually. The server itself would manage the on-disk storage of these files using a pool of disk space specified by the administrator. It would automatically backup some of the files to tape when this location gets filled up and, restore them later when needed for recovery.

In addition, the recovery infrastructure is being further simplified by eliminating many of the corner cases such as recovery through RESETLOG. The possibility of allowing the recovery of individual objects from a physical backup is also being investigated. Besides making the recovery process extremely easy, such simplifications will also improve the database availability by significantly reducing the recovery time.

A significant amount of development work is also going towards building intelligence into RMAN to help administrators plan the backup and recovery operations better. RMAN, in the future, may be able to predict the estimated completion time for backup or recovery operations. Administrators may also be able to seek advice from RMAN about their backup strategy based on a target recovery time. In other words, an administrator may be able to

ask RMAN to suggest a backup strategy in order to ensure that the recovery time never exceeds a given target. Similarly, it may be also possible to limit the impact of a backup operation on system performance by specifying a target completion time. RMAN may then automatically adjust its read rate to ensure that the backup operation is completed in the specified time window with as little impact on the system as possible.

PERFORMANCE TUNING

IDENTIFYING BOTTLENECKS

DBAs currently spend a significant amount of their time trying to identify the causes of performance problems. Diagnosing a poorly performing system is often considered a 'black art', surrounded by many myths and legends, few of them based on fact. An incorrect diagnosis often results in significant amounts of effort being spent to fix the 'problem', with little or no improvement in performance. For example, a DBA may spend weeks, sometimes months, reorganizing the disk subsystem because of a perceived IO problem, only to discover that the reorganization made little or no difference to the overall performance of the system since the bottleneck was due to unusually high parse rate.

A number of third party tuning tools are also available today but few of them are geared towards answering common questions such as *'Why is the system running poorly?'*, *'How can I make the biggest improvements in the system?'*, or *'Why is the system running worse today than it was last week?'*. Most products simply provide a graphical display of database statistics, leaving users to determine the root cause on their own by drilling through large amounts of data. Even with the best of tools, this can be a complex and extremely time consuming task.

Recognizing these problems, Oracle is currently developing solutions to automate the task of diagnosing performance problems. In the future, the database may be able to identify the performance bottlenecks proactively. The problem may then be either fixed automatically, if possible, or notified to the administrator along with a set of recommended solutions. A major objective of the automated monitoring process will be to identify, and eliminate from the tuning process early on, those parts of the system where tuning is unlikely to produce any substantial benefit. Automatic identification of the 'actual' bottlenecks in the system will greatly reduce the time required to analyze and resolve performance issues. This will also help direct manual tuning efforts towards the areas where maximum performance gains could be made.

Previous attempts and research on self-tuning databases have taken a bottom up approach with each component in the system continually trying to optimize its own operations. These have resulted in either point solutions for individual components, such as index advisors, or have been abandoned as too difficult because of complexity in the interaction between components. The methodology adopted by Oracle, therefore, seeks to perform an iterative top-down analysis of the system in order to focus the tuning activities on the sources of current bottlenecks. Such an approach will enable the database server to examine the bottlenecks in a holistic manner and, provide recommendations relating to all major system components including the operating system, the application, and the hardware.

SQL/APPLICATION TUNING

Application design issues are the most predominant cause of performance problems. All the tuning wisdom of developers, DBAs and system administrators cannot make up for the slowdown caused by architectural and design deficiencies of the application. While there is little that the database can do to fix such problems automatically, it may be able to identify them and recommend possible solutions. For example, the server may detect an unusually high rate of repeated logins from an application and recommend exploring the possibility of sharing pre-connected sessions.

In many cases, the performance degradation is caused by a few inefficient queries that consume most of the system resources. While Oracle's cost based optimizer (CBO) generates optimal execution plans most of the time, manual query tuning can lead to more efficient plans in some cases. This is particularly true in large DSS environments where execution of extremely complex queries, such as those involving joins across many tables or the ones using complex views with multiple layers of nested queries, is fairly common. Since missing or stale optimizer statistics is often the cause behind generation of sub-optimal execution plans, Oracle intends to automate the maintenance of such statistics. The database, in the future, will be able to create and refresh these statistics automatically eliminating yet another routine administrative task. The query optimization technology is also being enhanced with the goal of

creating a perfect optimizer – one that does not need any human assistance, never makes a mistake, and automatically adapts itself to changing workload. In the rare eventuality of the optimizer failing to make the right decision, a self-learning mechanism will ensure that such mistakes are never repeated. Such optimizer enhancements may completely obviate the need of any manual query tuning in the future.

DATABASE/INSTANCE TUNING

Effective utilization of system resources can have a large impact on overall database performance. Administrators therefore, continuously strive to tune the resource configuration parameters to maximize system performance and ensure the most efficient use of available resources. Simplifying the management of instance resources, therefore, is a central goal of Oracle's manageability efforts. Oracle9i Release 1 introduced the Automatic SQL Execution Memory Management feature to automate the management of the Program Global Area (PGA) memory. Using this feature, administrators just need to specify the maximum amount of PGA memory an instance can use. The server automatically distributes this memory among various active sessions in a manner that results in maximum performance gains. Going forward, Oracle is trying to develop similar self-tuning capabilities for the System Global Area (SGA) memory as well. Much of the groundwork required to achieve this goal has already been completed in Oracle9i. In Oracle9i, administrators are able to dynamically add to or remove memory from an active Oracle instance without shutting it down. Oracle9i also helps administrators in sizing the buffer cache optimally by predicting the "miss" rates for various sizes of the cache. Similar advisories for other memory components such as the shared pool and the PGA will be introduced in Oracle9i Release 2. These and other such advisories currently being developed will enable the Oracle server to completely automate the task of instance memory tuning in the future. The management of various memory components may also be integrated providing administrators with a single parameter to control the total memory consumed by a database instance.

The management of IO is another major administrative challenge. It becomes even more complex when using intelligent storage array products since the disk topology information for such devices is not easily available. Oracle9i Release 2 will have the capability to show a complete mapping of database files to intermediate layers of logical volumes manager and actual physical devices. This information will help administrators determine the exact location of a data block on the disk and easily identify any hot spots. In the future, the server may use this data, along with additional IO statistics, to automatically detect and alleviate IO bottlenecks. For example, if it is observed that the majority of IOs are concentrated on a small number of disks, it may be possible to automatically rearrange datafiles to spread IOs more uniformly.

Other aspects of database IO management are also being simplified. The self-tuning direct IO feature in Oracle9i Release 1 enabled the server to dynamically adjust the number of direct reads based on the nature of the query, ensuring the optimal use of available bandwidth. In the future, the server may also be able to automatically adjust its background IO rate according to workload in order to ensure minimal IO waits. The Recover Cost Estimator feature of Oracle9i Release 2 will help administrators make an informed trade off between the recovery time and the performance overhead caused by checkpoint writes. By predicting the number of physical writes for different values of the recovery target parameter (FAST_START_MTTR_TARGET), this feature will make it simple for DBAs to choose the recovery time judiciously.

An important objective of database performance management is to ensure that critical business operations are completed within a stipulated time. Using the Database Resource Manager, administrators can currently distribute CPU resources among database users and applications according to enterprise business objectives. In the future, this functionality may be extended to all system resources, including memory and IO bandwidth, giving administrators complete control over their deployment. The ability to partition memory and IO bandwidth, in addition to CPU, will enable administrators to provide predictable service to mission-critical applications with unprecedented consistency.

ADAPTIVE PERFORMANCE MANAGEMENT

The only constant is change. As factors such as workload and data volume change, the database performance can gradually deteriorate over time. Performance management therefore tends to be an iterative task requiring periodic resource and configuration adjustments. A key objective of Oracle's manageability strategy is to make the database server fully adaptive to such environmental changes. The future releases of the Oracle server may automatically capture the historical information related to workload and the operating environment. Further analysis of this data

may allow the server to detect changes and automatically initiate corrective actions, which at times may mean advising the administrator to acquire additional hardware resources. Also, analyzing the trend information related to periodic workload changes may enable the server to prepare for such changes in advance. For example, a database supporting a mixed environment, with workload primarily being OLTP during the daytime and batch processing at night, may be able to automatically shift memory from the SGA to PGA at the end of the day in order to ensure the availability of adequate SQL work area (sort, hash, bitmap merge etc.) memory for batch operations. By anticipating workload changes and adapting to them automatically, the server will ensure the most optimal use of available resources and guarantee the highest performance regardless of the nature of workload.

WHAT DOES IT MEAN TO YOU?

A completely self-managing database will radically change the way enterprise data centers are managed in the future. The automation of routine administrative tasks will enable DBAs to concentrate their efforts in managing the end-user experience rather than supporting technology. As the database becomes totally self-contained, the role of database administrators will evolve from managing the database to strategic management of business information stored in it. With an increased focus on company's business goals, database administrators in the future will play an increasingly crucial role in success of the enterprise.

Businesses will reap the benefits of increased administrative productivity in form of significantly reduced operational costs. At the same time, systems will perform better, offer more reliability, and automatically scale to millions of users – all at a fraction of today's costs. The “invisible” Oracle database will thus enable enterprises to become more profitable, provide better customer service, and emerge as market leaders.

The self-management capabilities of the Oracle server will also facilitate its use as an embedded database. Third party application developers will find it extremely easy to create applications that automatically manage the underlying Oracle database. The unrivaled performance, availability, and scalability of the embedded Oracle database will enable such applications to offer highest class mission-critical services without exposing database management complexities.

CONCLUSION

Modern enterprises are aggressively adopting new technology solutions to enhance their competitiveness and profitability. The system management costs, however, have been steadily rising leading to eroded profit margins. Oracle's war on complexity is the answer to this challenge. By undertaking a relentless campaign to simplify the management of database systems, Oracle is helping customers minimize administrative costs and free up corporate resources for those activities that truly differentiate a business from its competitors.