

New Features of Oracle8i and Oracle9i

TRACK DESCRIPTION

WatchIT.com™ brings together the best in IT. Our monthly programs feature technology experts, business leaders and industry analysts who offer a comprehensive perspective on the trends, issues and technologies important to today's IT professional. Understanding, advice, best practices: you get all of this and more from WatchIT.com™.

This program is part of our Technical Briefings track, which presents technology challenges, their solutions and the vendors that provide them.

INTRODUCTION

As you probably already know, the Oracle Database currently dominates the market and database software. Through the release of Oracle8, Oracle8i and now with Oracle9i, Oracle has evolved from a relatively simple relational database into one of the world's most sophisticated database solutions.

The challenge for IT for professionals is to ensure that they're able to leverage Oracle's powerful features to improve the productivity of your organization.

With the release of Oracle8i and Oracle9i, Oracle now offers a bewildering array of rich and complex new features. The ongoing challenge for IT managers is to sort through all of these features and understand which of them are the most appropriate to your own environment. An even more formidable challenge is in separating the marketing hype from the reality of the features. That's the main focus of this session.

By watching this program, you will get an in-depth technical understanding of how Oracle Database Software is involved and understand the new features that relate to e-commerce, data warehousing and very large databases.

As a DBA with more than 20 years' experience, I have learned, firsthand, which features are well received by practitioners and which ones have fallen by the wayside.

Welcome to the WatchIT.com™ Technical Briefings program: New Features of Oracle8i and Oracle9i.

I'm Donald K. Burleson, an independent Oracle database consultant. I specialize in creating database architectures for very large online databases and I have been fortunate enough to work with some of the world's most powerful and sophisticated Oracle systems.

As a leading database author, I have also written 12 books and published more than 100 articles. I currently serve as editor in chief of Oracle Internals, a leading Oracle database journal.

Recent books that I have published include:

- ~ The Best of Oracle Internals;
- ~ High-Performance Oracle Database Applications;

AGENDA

- ~ High Performance Tuning with STATSPACK; and
- ~ UNIX for Oracle DBAs.

I've worked with numerous Fortune 500 companies creating robust database architectures for their mission-critical systems.

Today the agenda includes:

- ~ A brief introduction of Oracle Database Software;
- ~ The new features of Oracle8i;
- ~ New features of Oracle9i;
- ~ Making the upgrade decision;
- ~ The impact on IT staffing and operations; and
- ~ Summary and recommendations.

AN INTRODUCTION TO ORACLE DATABASE SOFTWARE

Over the past 10 years, Oracle has emerged as the dominant database management system primarily because of its ability to run virtually every industry platform, all the way from mainframes to personal computers.

Because the Oracle software clearly commands the market for large database systems, IT managers everywhere are struggling to understand how they can leverage the Oracle technology to benefit their enterprises.

With the introduction of Oracle8, Oracle has made a commitment to start providing object-oriented technologies within the domain of its database engine. As Oracle8 evolved into Oracle8i, Oracle enhanced its software to provide Internet-ready API interfaces and other tools to allow Oracle to function in the Internet marketplace, especially in e-commerce systems.

More recently, with the introduction of Oracle9i, Oracle again reaffirmed its commitment to e-commerce vendors and Internet platforms, offering an entirely new suite of tools that are geared directly towards those who are using Oracle for high-volume, high throughput applications over the Web.

As a direct result of these enhancements, Oracle has gone from being a moderately complex relational database into one of the world's most sophisticated database software packages ever created.

It is not enough for the Oracle manager today (Q3, 2001) to be a generalist. Instead, he or she must understand all the intricacies and nuances of each of the many Oracle product suites. It's only through this detailed knowledge that an Oracle manager can appropriately leverage the Oracle technology for the maximum benefit to their own organization.

The Focus of the Oracle8i Release

The release of Oracle8i focused on improving the database engine to support Internet-based applications. First and foremost, Oracle made a commitment to support very large types of applications.

Whereas the Oracle7 database was very good for moderately sized applications – up to 100 GB – the Oracle8i software was designed for very large database systems that consume terabytes or even petabytes of information.

In order for Oracle to support this kind of huge scale, Oracle introduced several significant features relating to very large databases, mostly in the area of:

- ~ Optic partitioning;
- ~ Data warehouse tools such as Oracle Express; and
- ~ Database administration tools such as bitmapped indexes and star schema representation.

Originally, Oracle CEO Larry Ellison promised to rework Oracle8 into an object-oriented database. He later backed down because of a high volume of complaints from consumers who did not want to change from a relational architecture to an object-oriented system. Consequently, the Oracle8 database model is an object relational hybrid kind of tool.

Today's Oracle database remains relational, but significant object-oriented extensions have been added to the engine to allow for rudimentary object-oriented database functions.

One of my roles as an Oracle database expert is to determine which capabilities are important to users and which are not. In keeping with this goal, I will focus on the Oracle enhancements that are being adopted by the marketplace.

NEW FEATURES OF ORACLE8I

This discussion will be grouped into two areas and I am going to start by reviewing the base database improvements and then take a look at the new features of Oracle8i. I will present the new features in the order that I have seen them being embraced in the marketplace.

Oracle8i and Oracle9i have introduced improvements in several major areas, specifically, the performance and manageability of Oracle databases. These enhancements can greatly speed the execution of Oracle database queries and make the database run far more efficiently than ever before.

Optimizer Plan Stability

With the introduction of Oracle8i, an exciting new feature was introduced called optimizer plan stability, which is also known as stored outlines. Optimizer plan stability allows for the SQL tuning professional to store an outline for the execution of a SQL statement. This stored outline has the dual purpose of eliminating the parsing time required to generate the execution plan for an SQL statement and it also ensures that the same execution plan is always generated.

For data warehouse queries, the SQL parsing time for six way table joints could often take more than 30 minutes, and the use of stored outlines allow pre-tuned SQL statements to bypass parsing and begin immediate execution. Stored outlines are also a godsend for tuning application suites such as SAP and PeopleSoft where the SQL cannot be changed.

Cursor_Sharing: A New Initialization Parameter

Another exciting feature of Oracle8i is cursor_sharing. Cursor_sharing is a new initialization parameter in Oracle8i that is designed to help manage the clutter problems caused by nonshareable SQL statements.

When cursor_sharing is set to force (cursor_sharing=force) Oracle adds an extra layer of parsing that identifies the statement as an equivalent if the items differ only in the value of the literals hashing them to identical library cache objects. We will see that under the right circumstances this setting can help solve SQL performance problems relating to literal SQL within your library cache.

Setting cursor_sharing to force (cursor_sharing=force) may be worth considering if your system has any one of these following characteristics. They are:

- ~ A large number of statements in the shared pool that differ only by the values of their literals;
- ~ Response time that is slow due to a very high number of library cache misses. This can happen in hard parses in the library cache and is evident sometimes by latch contention.
- ~ In cases where the DBA has added cursor_sharing equal to force (cursor_sharing=force), this directive has worked very well in making SQL statements reusable, dramatically reducing the strain on the shared pool within Oracle. The downside is Oracle's warning that some SQL optimization plans will become sub-optimal, but this can easily be tuned by using optimizer plan stability.

Index-Organized Tables

Beginning in Oracle8, we have the ability to create index-organized tables, but unfortunately until Oracle8i they were not fully functional. In cases where an SQL query has many columns specified in the "where clause," it is possible for Oracle to service the query without ever touching the table itself. To do this, it is important that you understand that the Oracle index exists that has all of these fields available, to avoid a full table scan. This is called a fully concatenated index.

In such a case where the entire table participates in the index, the symbolic key information is stored entirely within the Oracle b-tree and the table itself becomes superfluous information. In other words, the index alone contains all of the information necessary to service any SQL query against that table.

Recognizing within these cases the table itself becomes superfluous, Oracle8i created the index organized tables whereby a table can exist with all of the information stored within the b-tree structure and avoid the overhead of actually having the table exist. There are numerous enhancements for index-organized tables intended to make them usable for very large databases as well as mission-critical applications.

Object-Oriented Extensions

The movement of Oracle8 toward object-orientation has resulted in some significant changes to the Oracle SQL syntax. As the Oracle DBA charged with tuning, it is important that we understand these extensions to Oracle SQL syntax and see how they affect performance.

The object-oriented extensions of Oracle SQL fall into four major areas.

- ~ The first is abstract datatypes. Oracle8 allows for the creation of abstract datatypes, which are sometimes called user-defined datatypes. These data types greatly simplify the Oracle's table structures and help you to create uniformity within your Oracle database.
- ~ We also see repeating columns within Oracle tables. These are known as non-first-normal form tables. Oracle8 now allows these non-first-normal form tables by adding repeating groups using Oracle's varying array or v-array data items within a single row, thereby representing a one-to-many relationship between data entities in a single data item.
- ~ We also see nested tables. Oracle8 introduced a new data structure concept whereby a column in a table has pointers to a nested table.
- ~ We also see using pointers to reference database rows. Oracle8 introduced a new concept whereby an object ID, which is commonly called an OID, can be used to directly grab an Oracle data row. This is a radical departure from the standard Oracle relational model, which mandates that we do not use pointers in order to navigate through the database.

However, pointers are a very necessary component of an object relational hybrid system. Unfortunately, the object-oriented extensions have not been widely accepted in the market, primarily because of the inherent complexity of object-oriented programming and the lack of availability of trained object-oriented personnel.

In addition to the enhancements to the core database engine, we also see a host of add-on products that improve the productivity of Oracle within selected areas.

Because of Oracle's commitment to delivering large-scale Oracle databases on the Web, the enhancements in Oracle8i provide several important new features in order to make databases easy to use over the Web.

The Oracle WebDB Product Tool

The most important of these tools is the Oracle WebDB product, which is an evolution of Oracle's Web server product from Oracle8. This is an easy-to-use Web development platform with a built-in API to communicate directly with the Oracle database from HTML. The WebDB tool allows developers to very quickly create HTML screens that contain Oracle data and also create HTML screens that are built dynamically, depending upon the contents from the Oracle database.

The Oracle WebDB product replaced the original Oracle WebServer product, which was introduced back in 1996. Oracle WebDB competes in the Web market space with other products such as Microsoft FrontPage, ColdFusion and with Oracles' own Oracle9i Internet Application Server product. The big distinction between WebDB and other third-party products is the ability to create native Oracle APIs, which we can use to build and extract dynamic information directly from HTML pages.

The WebDB product also allows for the easy creation of applications that will accept data on an HTML screen, transform the information and store that information inside the Oracle database.

In Oracle8i and on through Oracle9i, Oracle is also enhancing the WebDB product to allow ease-of-use when communicating with external portals, specifically with XML.

Oracle's New XML Interfaces

With Oracle's new XML interfaces, special Oracle listeners can intercept incoming XML messages from anywhere on the Internet, translate these messages according to a set of document type definitions or DTDs, and then build SQL transactions that can communicate directly with the Oracle database.

In the opposite direction, Oracle can respond to an XML request for data and build the Oracle data into an XML string. The string can then be transmitted across the Internet back to the requesting system.

Oracle claims that the WebDB environment is faster than the other professional tools. In my own experience, this is true for those who are not experts already in products such as FrontPage or ColdFusion. However, an IT organization's full-time Web developer staff are accustomed to using embedded APIs to communicate with databases to fetch and store information.

In summary, WebDB is a great tool for a neophyte who needs to create quick and robust Web applications that manipulate Oracle data. However, WebDB does not possess the same wealth of robust features that the generic Web development tools possess. Most full-time professionals would be very reluctant to abandon their regular Web development environment tools.

The Oracle Internet File System

According to Oracle, the Oracle Internet File System combines the power of Oracle8i with the ease-of-use of a file system. Internally, the Oracle Internet File System is a Java-based application that runs within the Oracle8i Java Virtual Machine, or JVM.

Oracle's Internet File System makes it possible to treat the databases as if they were simply a shared network drive. Users can store and retrieve files managed by the databases if they were using files managed by a file server. However, to date, there has not been a lot of interest in the Internet File System in the application development community, primarily because large applications do not have any need for their large databases to look like a file system.

The Oracle interMedia Product

The Oracle8i interMedia product allows the use of nontraditional data types within the Oracle database Web server application. These nontraditional data types may include large objects that could contain:

- ~ Wave files;
- ~ Images;
- ~ Videos; or
- ~ Any other graphical non-tabular types of data.

While this ability to store non-textual information is vital to Oracle as a company, the mainstay for most IT organizations continues to be text information. Well over 80% of Oracle shops have no need for storing images, audio, video, and the like.

Oracle's ConText Product

The Oracle ConText product is one of the most exciting new developments within the Oracle8i technology. Oracle ConText is essentially a concept-based search engine, which allows textual information to be very quickly indexed and retrieved by Web-based applications.

In today's marketplace, Oracle ConText competes with products such as Excalibur and NextPage. The idea behind these tools is that the user can highlight a paragraph anywhere on the screen, press a button labeled "show me more like this," and quickly see on point related information.

All concept-based word search engines are rated along two dimensions: precision and recall. They both use complex synonym expansion techniques and word stemming to get the best results. In practice, ConText's precision, recall and relevancy ranking functions are not as refined as other third-party products, primarily because Oracle ConText is a nascent entry into this market. Hence, while ConText is constantly being improved, many IT shops pursue other third-party vendors.

Oracle XML Solutions

With the expanding interest in e-commerce, many companies are struggling to find Internet-based methods for facilitating communications between e-commerce portals over the Internet. The extended markup language, commonly known as XML, has emerged as the dominant method for communicating between databases.

Oracle, as well as other vendors, is now marketing tools to assist in communications between Internet portals.

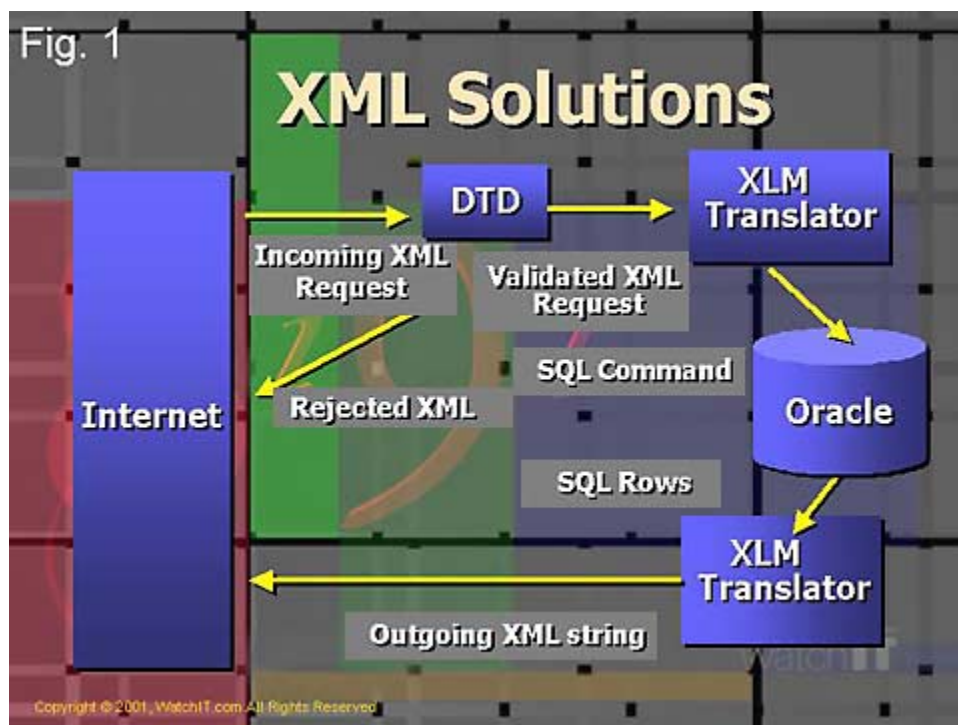
At a high level, XML defines a common protocol for communicating between databases. A document type definition, or DTD, is used to ensure that all the communications follow structured rules and that the Internet communications are achieved via standard HTTP calls that can be directed to special ports that listen for incoming XML transactions.

Once received, the transaction is decoded according to the DTD, and an SQL statement is generated to communicate directly with the Oracle database.

Responses follow the same rules. Outgoing data is translated into an XML string, and this XML command is sent back over the Internet to the requesting server.

Oracle, along with other vendors, is developing tools to aid in this XML communications explosion. The primary advantage to the Oracle tools is that they're tightly coupled with the Oracle database and they're offered free of charge with the Oracle database software.

Figure 1



Oracle Enters the Java Marketplace

Oracle has entered the Java market in keeping with their goal of allowing Oracle applications to execute on any platform. Because Java dynamically generates the executable code – it is interpretive in nature – Java can be executed on any hardware platform, making Java ideal for applications that must cross machine architectures.

However, most Oracle server environments do not have a need for a language that can run on multiple operating systems, since most Oracle servers are homogenous. Hence, while the Oracle and Java tools have been widely used for client-side applications, the use of Java Virtual Machine, or JVM, SQLJ, Java Stored Procedures, and Enterprise JavaBeans (EJB) have failed to gain widespread acceptance in the Oracle marketplace.

Again, this is because, primarily, the Oracle environment is always running a static server architecture and the interpretative nature of Java makes it too slow for most production applications. And the marketplace has been slow to adopt this new Java protocol primarily because of its slow speed.

Data Warehouses and Very Large Databases

Oracle8i and Oracle9i have introduced some new features that are directly targeted towards those IT shops with very large Oracle databases. These new features can improve the performance and manageability of very large Oracle databases, and also allow Oracle to support terabytes of online data. The Oracle8i and Oracle9i features relating to very large databases fall into several areas.

- ~ Materialized views;
- ~ Transportable tablespaces; and the use of
- ~ The Oracle Express product.

Oracle's Materialized Views Feature

Oracle has a special feature called materialized views that can greatly speed up data warehouse queries. In a materialized view, a summary table is created from a base table and all queries that perform a similar summarization against the base table will be transparently rewritten to reference this pre-built summary table.

These aggregations are re-computed on a schedule, normally during evening hours, and they serve to give the end users instantaneous response time to the summarization and aggregation queries.

A materialized view is a stored summary containing references to the pre-computed results. With the Oracle8i summary management feature, the Oracle database server automatically rewrites queries to use the summary tables rather than retrieving the data directly from the detailed tables by doing very expensive joins and aggregation operations. This query rewrite facility is totally transparent to the application, which is not even aware that the materialized views exist.

Transportable Tablespaces and Locally Managed Tablespaces

Oracle's transportable tablespaces are an outgrowth of the Oracle8i feature whereby tablespaces' metadata information is stored internally within the tablespace, instead of being stored inside the Oracle data dictionary.

Because all the metadata about a tablespace is now stored within the tablespace itself, it allows Oracle DBAs the ability to replicate a portion of a database by lifting and transporting the tablespace files from one server to another server anywhere in their network.

This transportable tablespace functionality dramatically improves the ability of Oracle administrators to manage read-only replicated systems. The tablespaces can be easily lifted and added to another existing database anywhere in the world.

Transportable tablespaces are not the same as Oracle parallel server systems, however, where multiple SGAs share a common tablespace. Rather, the tablespaces are data replicated into another remote database via a copy utility.

Of course, you must structure your database such that you have no cross-tablespace referential integrity constraints and the tablespace can exist as a finite entity.

In a data warehouse environment, where huge amounts of data flow to and from initial OLTP databases into your enterprise data warehouse, this feature presents an opportunity for faster and more innovative means of data movement.

The Oracle Express Product

The Oracle Express product was originally purchased by Oracle from IRI Corporation to provide multidimensional database capabilities for the summarization and aggregation of large amounts of factual information from a data warehouse.

The Oracle Express Engine allows end users to view the summarized information according to any number of salient dimensions, much like a pivot table does in an Excel spreadsheet.

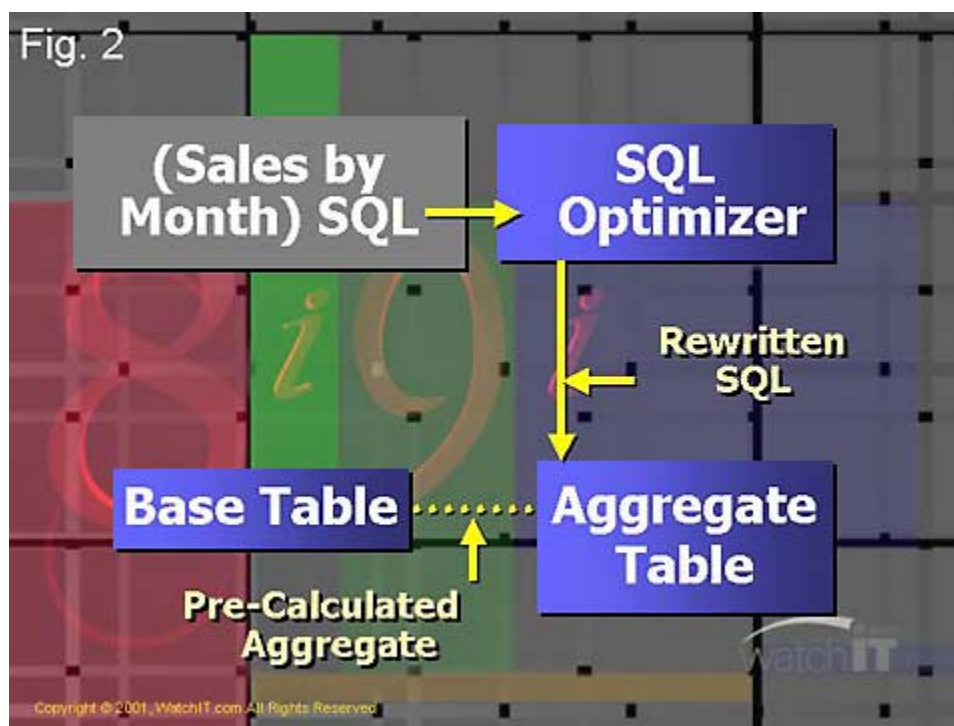
In a data warehouse, we have a:

- ~ Central fact table with items such as gross sales or total sales; and
- ~ Dimension tables that are used to summarize the data according to dimensions such as date, region or the type of product.

When Oracle detects an end user is summarizing or aggregating information, Oracle interfaces with the new materialized views feature to enable query rewrite transformations.

For example, if an Oracle session requested summing up of sales by month, the materialized views can detect that an aggregation has already been pre-calculated and go to the Oracle database to retrieve this aggregated information. This avoids the need for a very expensive full-table scan against the base tables that would need to happen if we were computing these averages in real time.

Figure 2



Oracle has a wealth of new features that improve security, both within the boundaries of the database and outside Oracle in the network.

Prior to Oracle8i, it was very difficult for the Oracle DBA to restrict access to specific rows within a table. However, there's always been the need to restrict access to individual rows in a table, especially in large corporations.

RADIUS:

A Client/Server Security Protocol

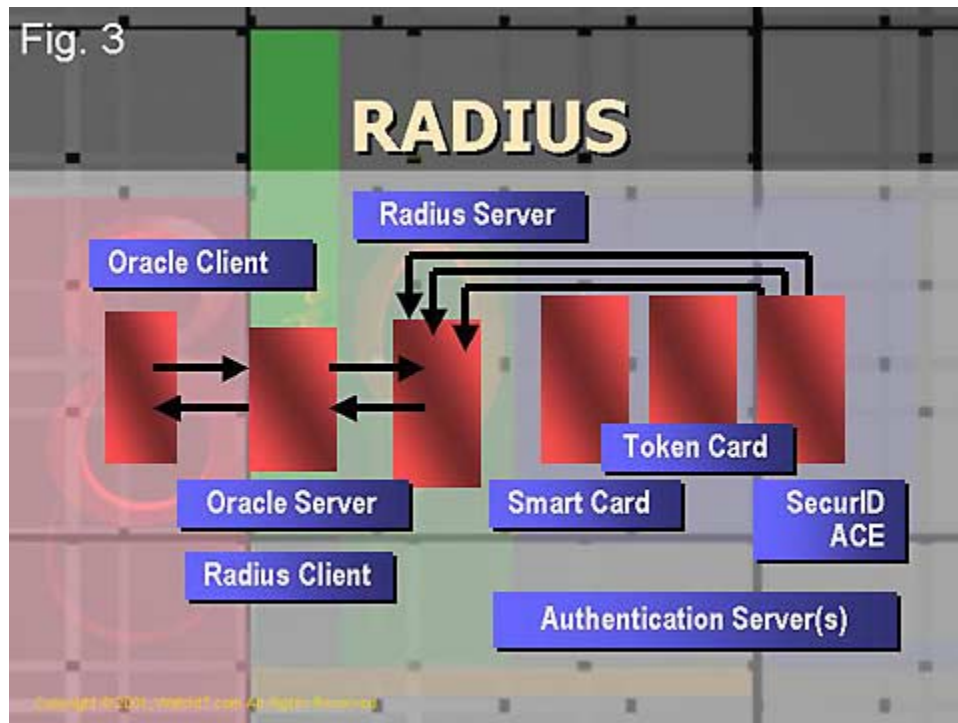
RADIUS, the Remote Authentication Dial-In User Service, is a client/server security protocol, most widely known for enabling remote authentication and access to Oracle databases. Oracle Advanced Security Features in Oracle8i uses this industry standard in a database, client/server network environment.

With RADIUS, you can enable the network to use any authentication method that supports the RADIUS standard, including token ring cards, smart cards or anything else, simply by installing and configuring the Oracle RADIUS feature.

Moreover, when you use RADIUS, you can change the authentication method without modifying either the Oracle client or the Oracle database server.

From an end user's perspective, the entire authentication process takes place seamlessly and transparently. When the user seeks access to an Oracle database server, the Oracle server, acting as the RADIUS client, notifies the RADIUS server, which in turn acts as an authentication server to verify the transaction.

Figure 3



NEW FEATURES OF ORACLE9I

Now, let's take a look at the new 9i enhancements to Oracle.

In Oracle9i, Oracle has made significant improvements in several major areas of the base database product. The most important of these are related to improved speed of processing, as well as the incorporation of Apache code into the Oracle Internet Application server.

As you may know, Oracle has a long history of introducing new releases with a huge amount of fanfare. As an Oracle professional, it takes me many hours to translate the 9i marketing hyperbole into a list of real features. Hopefully, this overview will shed some light on what these new features will mean for you.

Let's begin with the basic database enhancements that I feel are the most important in Oracle9i.

In-Place Table Reorganization

In Oracle8i, the DBA was always forced to take the database tables offline in order to reorganize the tables.

As you might know, reorganizations are used to:

- ~ Defragment Oracle rows;
- ~ Remove chained rows; as well as
- ~ Coalesce the freelists in the table segment headers.

This used to involve a manual export and import operation, or the use of Oracle's Create Table as Select, commonly known as CTAS, to copy the table for reorganization.

In Oracle9i, Oracle has incorporated the CTAS utility into a new facility that copies a table into a temporary area and then automatically reapplies all indexes and referential integrity constraints to the table.

Recovery From Online Redo Logs

Oracle9i also supports recovery by reading the online redo logs. This imitates a product by Quest, known as SharePlex, and allows the Oracle9i system to implement standby databases with almost instant synchronization capabilities.

Oracle9i Support for Databases With Multiple Block Sizes

Another exciting new feature is that Oracle9i also supports databases with multiple block sizes. This has been a feature available in other nonrelational databases for decades and it's exciting because it allows administrators to configure corresponding buffer caches for each alternative block size within their Oracle database.

This allows Oracle DBAs to segregate large row tables into files with large block sizes while keeping small row size tables in files with smaller block sizes. This can have a huge impact on the disk I/O system and minimize I/O-related performance bottlenecks.

The SGA: Dynamic System Global Area

With Oracle9i the DBA can dynamically change the memory allocation within the system global area, commonly called the SGA. This is a remarkable breakthrough for Oracle because now the DBA can automatically monitor Oracle, as well as the database server, and adjust the size of the SGA memory region depending upon the current processing that's going on within the Oracle9i database.

Oracle9i will create a foundation for a self-tuning database instance, and this is a very exciting new development within Oracle.

The Oracle9i Self-Tuning Database

With the dynamic SGA features, Oracle 9i has created a foundation for a self-tuning database and this is a very exciting new development within Oracle.

This memory management feature is another area that has been given significant attention in Oracle9i.

Traditionally, Oracle DBAs have needed to shut down the database instance in order to grow or shrink specific SGA components. With Oracle9i and the dynamic memory management feature, we can resize:

- ~ The data buffer cache;
- ~ The shared pool; as well as
- ~ The sort area size or the large pool, directly with online commands keeping our system available 24 by 7.

Furthermore, Oracle9i provides for transparent management of working memory within the SQL execution by self-tuning the initialization runtime parameters controlling the allocation of private memory.

Now, let's move away from the base 9i features and take a look at some of the new 9i features and tools.

Beginning with Oracle8, Oracle made a commitment to support data warehouse types of applications. In fact, the major features in Oracle8 were strongly centered on database scalability and Oracle's ability to store many terabytes of information for the user community. Oracle is continuing this commitment to online analytical processing, commonly known as OLAP, and data warehouse tools in Oracle9i.

Let's take a look at some of these new data warehouse features.

Oracle9i Data Warehouse Features

Oracle9i SQL has changed significantly over the past four years. Oracle Corporation has been struggling to find a balance between making SQL adhere to the ANSI standard and providing robust built-in functions that are necessary in order to improve Oracle SQL performance.

Oracle9i has deliberately chosen to vary from the ANSI standard from SQL and provide a wealth of built-in functions that can improve the performance of the Oracle database.

Oracle has also improved the responsiveness of SQL by providing Java-based interfaces that go directly into the OLAP engine for data warehousing. This makes SQL much more accessible to external portals from within the Internet world.

These Oracle9i SQL enhancements include the LEAD and LAG functions. These SQL extensions enable complex time-series analysis of data warehouse information and allows for the easy aggregation of rolling time periods, which is commonly required in a large data warehouse.

The Oracle9i SQL also has new enhancements to ranking functions. Oracle9i SQL includes ranking functions that provide support for common OLAP rankings, such as:

- ~ Show me the top 10;
- ~ The bottom 10;
- ~ The top 10%; or
- ~ The bottom 20%.

Oracle9i SQL Advanced Grouping Functions

Oracle9i SQL also includes advanced grouping functions. Oracle9i now provides categorization functions that group values into buckets, such as age groups or income brackets.

Oracle9i also adds some new statistical functions. They include:

- ~ Support for correlation analysis;
- ~ Standard deviation;
- ~ Linear regressions using single, double, triple exponential smoothing as well as some of the least squared; and other statistical calculations that are commonly used within a data warehouse.

Oracle9i OLAP: Extract Transform and Load – ETL

In addition to the SQL enhancements, Oracle9i has improved its OLAP engine and offers tools that simplify the complex task of loading Oracle data warehouses.

The first of these is the extract, transform and load, commonly known as the ETL tool. ETL functionality has been a long-awaited feature in Oracle data warehousing.

One of the most confounding problems for Oracle DBAs managing a data warehouse is providing clean and accurate summarized data from the operational data stores. Getting it into the warehouse in a timely fashion has always been a challenge.

Prior to the introduction of ETL software, the extracting and loading of Oracle information tended to be very time consuming and cumbersome. Customized programs had to be written to extract the data from the operational data store.

Transformation and summarization was largely done manually. Loading the information into the target data warehouse would often require the dedicated efforts of a full-time DBA. Many of the ETL functions in Oracle9i were integrated through the acquisition of Carleton Software for data

warehousing tools. The purpose of the ETL software is to assist in the loading of the data from the operational data stores into the data warehouse.

The ETL offering includes the following functions:

- ~ The first is extract. The extract phase of the ETL software consists of a set of Oracle SQL queries that pulls the detailed information from the operational data store.
- ~ The next phase is transformation. The transform phase generally involves aggregating highly detailed data that was taken from the operational database.

Oracle9i claims to have tools that will assist in the automatic aggregation of information so that it can be summarized for easy loading into a data warehouse engine. The information-transform phase sometimes involves collecting external metadata and cleansing the operational data.

- ~ The load functions of the ETL software relates to the actual loading of the information after it's been extracted, transformed and cleansed and is ready to go into the data warehouse.

Prior to ETL, the Oracle DBA would've had to have written a customized SQL loader program in order to handle this task.

Oracle9i Personalization Technology Feature: The Recommendation Engine

E-commerce sites have been demanding tools that would allow the Web site to survey consumer behaviors and make product recommendations based upon the customer's prior searching, viewing and purchasing behavior.

Oracle9i introduces a personalization technology feature called a recommendation engine that simplifies the tracking of consumer behavior by collecting individual information for all consumers who visit the Oracle Web site. This data takes the form of Web server statistics.

Customers are demanding custom content on the Web, and looking at Amazon is a great example of a recommendation engine technology. The data from a recommendation engine can be used to target advertising or to customize Web page content.

Just as the Amazon home page displays products for you based upon your prior product viewing, the recommendation engine keeps track of all page visits and the duration of time spent viewing each page.

Obviously, the recommendation engine is aimed at retail e-commerce systems where consumers are actually signing on and surveying a broad range of products and services. With the recommendation engine, Oracle is attempting to add business intelligence to its e-commerce solution software.

The personalization feature will allow the data warehouse engine to automatically survey past usage trends and make recommendations for purchases based upon your end users' prior behavior.

The Improved Oracle Parallel Server – OPS

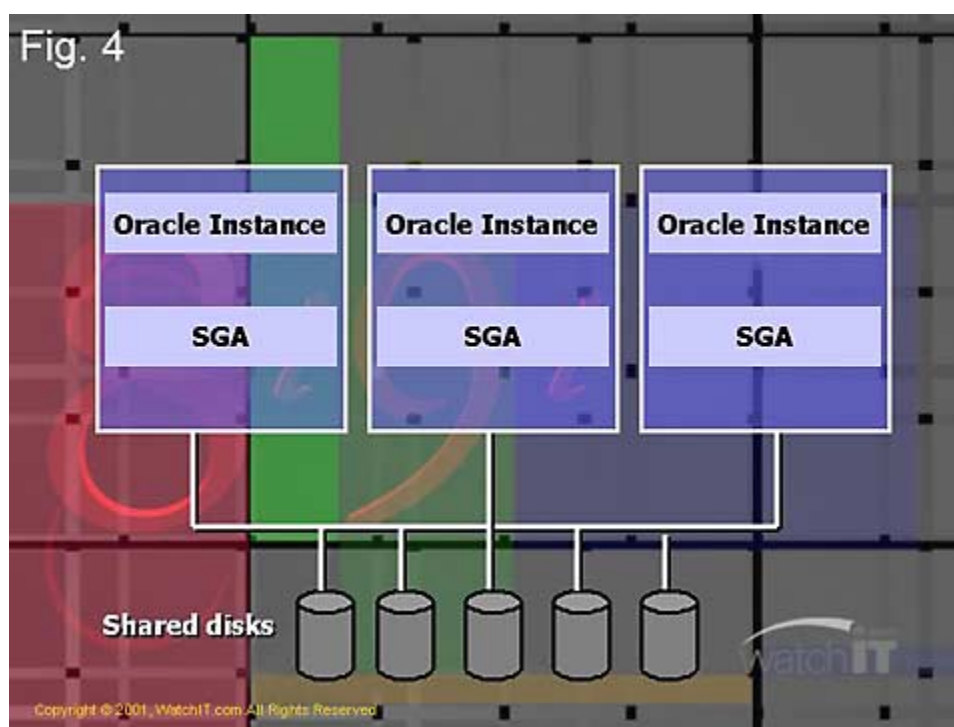
Oracle Parallel Server, or OPS, has been around for years.

Now, Oracle has taken the general idea of OPS and improved upon it, offering a new product, which they call Cache Fusion and Real Application Clusters, or RAC for short.

The term "cache fusion" describes the ability of a RAC database to allow multiple SGAs to share data buffer storage between the instances while still accessing a single Oracle database.

A RAC configuration requires that the instances be able to share a common data buffer cache. The concept of cache fusion is essentially advertising the fact that the Oracle Integrated Distributed Lock Manager, or IDLM, has been changed to allow for a common storage area for the data buffers between instances.

Figure 4



This greatly improves the manageability of parallel server systems, and makes Oracle9i fault tolerant, since the failure on one node will not result in a complete system crash.

Oracle9i combines this RAC failover feature with their transparent application failover or TAF function, which allows Oracle sessions that get disconnected from one server to quickly reestablish their session on another RAC node.

Java and XML Support Enhancements

Now let's take a look at the Java and XML support enhancements in Oracle9i.

The enhancements for Java and XML support services on Oracle9i further demonstrate Oracle's focus on the Internet database. Without going into a great level of detail about JavaBeans and XML communications, these tools enable distributed Oracle databases on the Web to communicate with each other using a standard interface protocol.

Oracle has made a commitment to extending its Java class structures and to provide easy-to-use XML interfaces in Oracle9i. Whereas a traditional Oracle database requires you to call information using Net8 database links, a database can now communicate with any other database on the Web simply by using XML. All that's needed to establish communications is the:

- ~ IP address;
- ~ The port number; and
- ~ The XML protocols for communicating with the remote database.

This is a very exciting new feature of Oracle9i and promises to make database communications of e-commerce portal systems a commonplace activity.

Oracle9i Data Mining Tools

The Oracle9i data warehouse suite also offers new data mining tools. As you might know, data mining is a general term for programs that examine large volumes of data and computes statistically significant correlations. The data-mining product that Oracle9i offers is an enhancement of the Darwin data-mining product, which Oracle purchased in 1999.

You probably already know that data mining is the most advanced of all the data warehouse facilities. Even in the simplest forms, data mining tends to be quite analytical and time consuming against the target database. Many data mining tools also have a tremendous amount of overhead because of their need to survey billions of bytes of information in an attempt to tease out statistically significant correlations between information in the Oracle data warehouse.

It remains to be seen how well the Oracle data warehouse community will accept this new data-mining tool. Oracle9i users may instead choose one of the many third-party data warehouse tools that can be easily attached to the Oracle database. These include established tools such as:

- ~ SAS;
- ~ SPSS;
- ~ SGI;
- ~ As well as other nascent companies developing data mining tools to go against the Oracle data warehouse.

THE IMPACT ON IT STAFFING AND OPERATIONS

Now, let's talk about the staffing impact on your IT department and how it relates to new releases of Oracle.

As we mentioned, Oracle has evolved into one of the world's most complex database management systems, and it takes years of training to master all the complexity of the Oracle database.

Just as the medical profession is segmented into specialty areas because of its complexity, Oracle DBAs are also segmenting themselves into specialty areas, targeting Oracle DBA niche markets such as:

- ~ Data warehousing;
- ~ Performance tuning; and
- ~ Oracle applications support.

It is very hard to find an experienced Oracle DBA who is a general practitioner. In addition, most Oracle DBAs are compensated at the same level as corporate vice presidents, and some senior DBAs in high salary areas exceed \$150,000 dollars per year.

Companies are struggling to attract DBA talent with enough experience to ensure continuous availability of their databases. The downtime for many of these production databases is often measured in tens of thousands of dollars per minute, and it is very risky for an IT manager to entrust these databases to beginners.

Approaches to the DBA Dilemma: Growing Your Own DBA

Since many companies have been unable to attract DBA talent, several approaches have been tried to fill the market need.

One alternative is to grow your own DBA. Some companies undertake to educate a promising junior employee in database administration.

For example, the cost of a full set of DBA training by Oracle Corporation costs about \$70,000 dollars, and many companies are successful in training these junior DBAs. Unfortunately, most companies lose these new DBAs as soon as they're trained because their HR departments are unable to increase their salary according to the market rate.

It is very difficult for many HR departments to justify the salary of the Oracle DBA at the same level of a corporate vice president, and it is impossible for many companies to raise the salary of an employee from, say, \$40,000 dollars a year to \$120,000 dollars per year. Hence, most of these DBAs leave their corporation soon after completing their training.

Approaches to the DBA Dilemma: Sponsoring a Foreign Professional

Another approach is to sponsor an alien DBA. There are major universities in India that produce thousands of Oracle DBA professionals every year targeted at the American marketplace. By agreeing to sponsor an alien DBA under his H1B visa, the employer is ensured that the individual will be available for several years.

Unfortunately, there are several serious problems with some of these employees. Some foreign universities do not pre-test applicants for their ability to perform DBA duties, nor do they weed out those students who do not show a natural predisposition to computing activities.

Approaches to the DBA Dilemma: Remote DBA Services

A third alternative is to pursue remote DBA services. Many companies are abandoning full-time employee DBAs in favor of remote Oracle support. These outside remote DBA companies leverage on a shared skill pool, and offer remote DBA support far cheaper than having your own expensive full-time Oracle DBA.

These services hire a pool of highly skilled Oracle DBAs and place them in response centers where they can constantly monitor their databases. There are currently more than 100 companies offering remote DBA services.

Training Your Oracle DBAs in New Concepts and Languages

Now, let's take a look at another area that might impact the IT department.

Additional time will be required for Oracle DBAs and developers to learn and embrace the new Oracle8i and 9i technology. This is especially important in the development areas where developers are required to implement strange new concepts and languages such as Java, XML, Apache and Kerberos. This will involve significant time off to attend the Oracle courses, and also significant expense for the training. Counting travel expenses, each week of Oracle training can easily cost \$4,000 dollars.

We also see an impact on IT roles with the rollout of Oracle8i and Oracle9i. Roles are going to change for the DBA because the DBA scope of responsibility must be increased into new database areas. In addition to current responsibilities, the DBA must also master:

- ~ Object-oriented features;
- ~ Java;
- ~ Kerberos;
- ~ XML; as well as;
- ~ The new enhancements and product tools such as Real Application Clusters; and
- ~ The new DBA management tools.

Roles for the developers will also change because of the new technologies that interface with Oracle. These again include:

- ~ Java;
- ~ XML;
- ~ Apache;
- ~ C++; and
- ~ A host of new and different coding techniques.

Over the long term, most conservative IT managers do not rush headfirst into cutting-edge technology without a pressing functional need. Hence, most IT managers only adopt the new Oracle8i and 9i features only when mandated by their own vendor and business requirements. Nobody wants their production system to be the first test site for new Oracle features, and the prudent IT managers will always wait until all of the bugs have been worked out.

SUMMARY AND RECOMMENDATIONS

Now let's take a look at some recommendations for when to adopt Oracle8i or Oracle9i.

While Oracle has introduced a plethora of new features and enhancements, most conservative IT shops are adopting a wait-and-see attitude and most companies will wait until the second release of Oracle9i before proceeding with the product. Remember, business requirements drive the technology and many of the new Oracle features are not utilized until there is a legitimate business need for this feature.

A notable exception to this is the base DBA features. It is perfectly safe to adopt some of the internal performance features, such as:

- ~ In-place table reorganizations;
- ~ Function-based indexes; and
- ~ Other database features that are relatively inert.

In summary, when embracing the new Oracle9i technologies, you must be very careful to ensure adequate staff training and to ensure that you fully understand how your business needs mesh with the new Oracle offerings.

FOR ADDITIONAL INFORMATION

Thank you for joining us for the WatchIT.com™ Technical Briefings program New Features of Oracle8i and Oracle9i. I am Don Burleson. If you have any questions or comments regarding this program, please feel free to e-mail me at: experts@watchit.com