Database Systems Management and Oracle8

C. Gregory Doherty
Oracle Corporation
gdoherty@oracle.com

1. ABSTRACT

Oracle's corporate mission is to enable the Information Age through *network computing*, a vision of broader access to information for all and the empowerment and increased productivity that can result. The technology implications of the network computing vision are ubiquitous access via low-cost appliances to smaller numbers of larger databases, accessed via professionally managed networks compliant with open internetworking protocols. The latest release of the Oracle data server, Oracle8, provides new technology for management of very large databases containing rich and user-defined data types, and is continuing to evolve to make it economically beneficial to store all forms of digital information in a database.

2. INTRODUCTION

The dictionary defines manageability as "the capability of being managed", and managed as "to handle or direct with a degree of skill as to make and keep compliant, to treat with care, to exercise executive, administrative, and supervisory direction". These are relevant words relative to data management, and bring to mind issues such as good vs. bad data models, application design for performance and scalability, operational administrative tasks to ensure data validity, recoverability, and high availability. Addressing these issues requires knowledgeable and skilled staff, both to design applications and to operate the databases on which they are built.

Oracle's view of the evolution of information systems is that it is strongly influenced by commercial economic factors. Today, more than ever, an economic inversion of the traditional relationship between technology and technology workers is evident: hardware and software is cheap and getting cheaper. Skilled IT professionals who can mold these technology components into useful systems are expensive and getting more expensive. Concurrently, the demand for useful information systems is exploding due to three primary factors: decreasing technology costs, rapid expansion of the Internet, and the requirement to make IT a vital, real-time component of the conduct of global commerce.

In a survey of over 500 companies, 82% said that they had plans to expand their IT staff, yet 68% cited a lack of skilled and trained workers as a barrier to that growth. The study also observed that the number of US graduates in computer science in 1994 fell by 43% (24,200 total), compared to 1986 (42,195 total) helping to balloon the costs associated with acquiring and maintaining an IT staff. [2]

In a survey of 700 organizations worldwide with at least 50GB of data under management, 94% report that their data centers are taking control of critical servers, an increase of 13% over the survey of the prior year. Over 50% of those surveyed cited improved information management, faster data access in multi-platform environments, faster response to business opportunities, and improved use of IT staff and resources as reasons for this consolidation trend. Over 50% of the respondents also cited poor customer service, lower business productivity, lost revenues, reduced competitiveness, and difficulties with mergers and acquisitions as some of the problems attributable to continued dispersed data. [1]

Databases will proliferate to support new applications, however, the trend to have databases consolidated into managed environments, and for each database to be larger as a means of amplifying limited numbers of IT staff seems to be a strong trend. Observing this economic trend, Oracle has focused its development efforts in Oracle8 on two primary areas: supporting collections of very large databases in an internetworked environment, and reducing the "impedance mismatch" between application design and traditional relational databases.

3. Intra-server Management (The Tools)

Oracle8 represents another step in the Oracle data server evolution and constitutes many hundreds of person years of development effort. Besides the portability of the Oracle server, allowing it to run on computing platforms of all architectures and sizes suitable for a given application, some of its key new capabilities to help address the aforementioned trends include the following key manageability tools:

3.1 Large Database Support

Our customers have been presenting requirements for individual databases in the terabytes or ten's of terabytes size, so a lot of work was done to increase the raw storage capability of the server, and concurrently to provide the management tools required to "divide and conquer" large data. We increased server address space to 8PB and allowed databases to be comprised of as many as 64K files, with each file being tens of GB in size. This was done with no requirement for reformatting of existing Oracle databases. Tablespaces organize database files into related groups, and Oracle8 allows logical objects (tables, indexes, etc.) to span multiple tablespaces. Data partitioning of tables and indexes allows logical placement of rows into up to 64K related partitions. The design of data partitioning placed particular emphasis on using partitioning to manage large data sets in manageable pieces, as distinct from systems that partition solely to achieve parallelism.

SIGMOD '98 Seattle, WA, USA

The query executor and optimizer have been improved to transparently distribute individual query workload over multiple processors in an SMP system as well as processors in a clustered system, and to fully exploit all optimizations made possible by logical data partitioning and new access methods such as bitmap indexes, found often in our customers' star schema data warehouse applications.

Various bulk management parallelized operations are provided for loading, moving, and removing data, as well as indexing, reorganization, and repartitioning, and parallel DML. These are useful as many Oracle databases are deployed to take bulk data feeds from legacy environments or proprietary devices, e.g. phone switches.

A new database cache design improves the ability of the server to work with large (10's of GB) caches and provides significant new algorithms for handling checkpointing of such large caches. The server also makes judicious use of ample RAM to perform optimization of sorting, indexing, and processing of hash joins.

A new large-object attribute type is introduced which allows for data objects up to 4GB in size to be manipulated piece-wise using a filesystem metaphor.

3.2 Large Client Populations

Large enterprise-wide client populations are now very common, and handling them in Oracle8 received a great deal of development effort.

Multiple client applications with low database access duty-cycles, common in many industries, can share a single virtual circuit to the database server, saving on client system cost and overcoming common OS network layer limitations. We also added the ability to deploy two-tier client applications in a three-tier multiplexed architecture, again sharing OS networking resources, without changing the applications.

DBA-defined role-based security mechanisms allow arbitrarily complex security enforcement for large numbers of named users across a significant number of grantable data access and administrative capabilities. We also added support for external authentication mechanisms, external directory services, and single sign-on in an internetworked security environment, commonly required in large commercial and government installations.

In the case of a server host failure in a clustered, shared database configuration, we created the ability to automatically reconnect client applications and restore query state to a surviving host serving the same database, saving clients the "reboot" time.

3.3 Database Interoperability

Oracle servers can distribute query and transaction workload amongst databases with transparent 2PC for multi-site DML. Oracle8 adds to the N-way multi-master replication and snapshot capability the support for logically defined data subsets, allowing for replication of only required data to remote sites, as defined by an application. This enables high availability, and is a foundation for disconnected client operation using our client database server, OracleLite. Oracle8 also debuts a new architecture for interacting with heterogeneous database systems and data sources, transparently distributing query and transaction workload amongst other data management systems. This is key for almost all enterprises that have a mix of hosts and legacy applications that must interact.

3.4 Application Development Productivity

To allow object-oriented applications low-impedance access to existing relational data we created "object views" that allow existing relational tables to be viewed, related and updated as object types.

We also created a new type system that allows new types to be defined and used either as attributes of tables or to implement tables of user-defined types. This functionality also includes all the programmatic interfaces to allow for client-side object caching, and object navigation, with minimized client-server messaging traffic, and with same consistent-read, snapshot-based concurrency control and isolation semantics of the Oracle server.

We debuted the ability to extend the functionality of the database server through "data cartridges", modules of application or user-defined code that are called out to securely from within the database server environment. Some specific cartridges implemented by Oracle include spatial data, text, and time series data.

4. Extra-server Management (The Toolbox)

A key technology for the management environment of Oracle database servers and their accompanying tools is manifest in the Oracle Enterprise Manager product. Enterprise Manager provides a Java and CORBA/IIOP-based client browser interface to a wide variety of functions that help to manage Oracle databases. The Internet-based architecture of Enterprise Manager allows for remote administration and monitoring from any browser-like client device, and includes a layer that implements management objects for one or more databases that can be local or distributed anywhere in the world. This architecture is important in that it allows an enterprise to have multiple databases, perhaps in many locations, but to be able to operate and administer them from one location to maximize IT staff productivity.

Some of the key functions that Enterprise Manager provides via its Internet-accessible management objects include database-driven automated backup & recovery, job scheduling, performance monitoring, diagnostic monitors, schema change management, tablespace management and reorganization tools, security, etc. Another key attribute of Enterprise Manager is its ability to perform actions common to multiple databases with one invocation. This is particularly useful when performing common administrative tasks like managing named users and roles, or performing schema changes across multiple databases.

Enterprise Manager is the framework and platform through which the extensive manageability tools in the Oracle8 server are presented. Administrators can use the tools and the toolbox in a way that amplifies their productivity, a key requirement in a time of IT staff shortage, proliferating databases, and larger more critical databases.

5. REFERENCES

[1] Find/SVP Strat. Res. Division, Managing Information Across the Enterprise: Closing the Information Gap, June 1997.

[2] Information Technology Association of America/Virginia Tech. Study, January 1998.