



# Craig S. Mullins & Associates, Inc.

*Database Performance Management*

[Return to Home Page](#)

May 1998

## **Computing : News & Review**

### The Most Important Thing is Performance

*By Craig S. Mullins*

Databases. No matter how feature-rich and complex the DBMS becomes the thing that is most troublesome for those of us who use them is performance. If the data is not readily available in an optimal manner, application usage will decline and the competitive advantages supposedly delivered by the application will not accrue.

Therefore, planning for database performance management is a crucial component of any application implementation. Without a plan for monitoring performance and tuning the database and the associated SQL, performance degradation is sure to occur. A complete performance management plan will include tools to help monitor application performance

and tune the database and SQL.

The most likely culprit for most performance problems is the SQL in the application code. In fact, industry consensus indicates that 75% to 80% of all database performance problems can be traced back to poorly coded SQL. This does not mean that the SQL in applications is bad to begin with. In fact, an application may be 100% tuned for rapid relational access when it first moves into production, but over time experiences performance degradation. This can occur for many reasons such as database growth, new data access patterns, changes in the business, etc.

Of course, the SQL can be just plain bad to begin with, too. Any number of problems can cause poorly performing SQL including:

- table scans
- lack of appropriate indexes (columns indexed)
- improper indexing choices (hash, clustering, uniqueness, etc.)
- not using appropriate indexes
- optimizer statistics not up to date
- joining tables in a sub-optimal order
- application joins instead of (usually) more efficient SQL joins
- improper join method (nested loop, merge scan, etc.)
- efficient SQL inside of inefficient application code (loops)

- inefficient subquery formulation (exists, not exists, etc.)
- unnecessary sorting (group by, order by, union)

Of course, there are other factors that can negatively impact database performance. It is wise to periodically check the overall performance of the database instance and the operating system it works on. Some quick items to check include:

- memory allocation (buffer/cache for data, SQL, etc.)
- logging options (cache, size, Oracle rollback segments, etc.)
- I/O (separation of tables and indexes on disk, database size, etc.)
- overall application and database workload on the server
- database schema definition

## **Database Performance Tuning Tools**

Tools are required in order to effectively manage database performance. There are many third-party tools that can help to effectively manage the performance of database applications. Tools that enable DBAs to tune databases fall into two major categories: performance management and fast utilities. Many different types of performance management tools exist that fall into one or more of the following categories:

- **performance monitors** - enable DBAs and performance analysts to gauge the performance of applications accessing databases in one (or more) of three ways: real time, near time (intervals), and/or based on historical trends. The more advanced performance monitors are agent-based.
- **performance estimation tools** - perform predictive performance estimation for entire programs and SQL statements based upon access paths, operating environment, and a rules or inference engine.
- **capacity planning tools** - enables DBAs to tune environment by analyzing the current environment and database design and performing "what-if" scenarios on both (environment and design).
- **SQL analysis & tuning tools** - provide graphical and/or textual descriptions of query access paths as determined by the relational optimizer. These tools can execute against single SQL statements or entire programs.
- **advisory tools** - augment SQL analysis & tuning tools by providing a knowledge base of tuning tips so suggestions can be made as to how SQL can be re-formulated for optimal performance. Advanced tools may automatically change the SQL (on request) based on the coding tips in the knowledge base.
- **system analysis & tuning tools** - enable the DBA to view and change database and system

parameters using a graphical interface (e.g. cache and/or buffer pool tuning, log sizing, etc.).

Furthermore, in the utility category, several tools can be used to tune databases:

- **reorganization tools** - databases can cause performance problems due to their internal organization (fragmentation, row ordering, storage allocation, etc.). Reorganization tools automate the process of rebuilding optimally organized databases.
- **compression tools** - enable DBAs to minimize the amount of disk used by databases thereby reducing overall disk utilization and, possibly, reducing elapsed query/program execution time because fewer I/Os may be required. (Caution: compression tools can also increase CPU consumption due to the overhead of the compress/decompress algorithms used).
- **sorting tools** - can be used in two ways: 1) sorting data prior loading databases to ensure that rows will be in a predetermined sequence, and; 2) in place of "order by/group by" SQL. Retrieving rows from a relational database may or may not be more efficient using SQL and "order by" versus SQL alone followed by a standalone sort of the SQL results set.

Frequently, each of the tools will need to be used in conjunction with one another. These tools work best if they are integrated and accessible from a central

management console. This enables DBA to perform core performance-oriented and database administration tasks from a single platform.

Many DBMS vendors provide solutions to manage their databases only; for example, Oracle provides Oracle Enterprise Manager and Sybase provides SQL Central for this purpose. Third party vendors provide more robust options that act across heterogeneous environments (e.g. multiple different database servers; DB2, Informix, Ingres, Oracle, Sybase, MS SQL Server, etc., running on multiple different operating systems).

In general it is wise to use the DBMS vendor solution only if your shop has one DBMS vendor. Organizations with multiple DBMS engines running across multiple operating systems should investigate the third party tool vendors.

### **Synopsis**

Applications that access relational databases are only as good as the performance they achieve. The wise organization will implement a comprehensive performance monitoring, tuning, and management environment consisting of policies, procedures, and integrated performance management tools and utilities.

From [\*Computing News and Review\*](#), May 1998.

© 1999 Mullins Consulting, Inc. All rights reserved.  
[Home](#). Phone: 281-494-6153 Fax: 281-491-0637