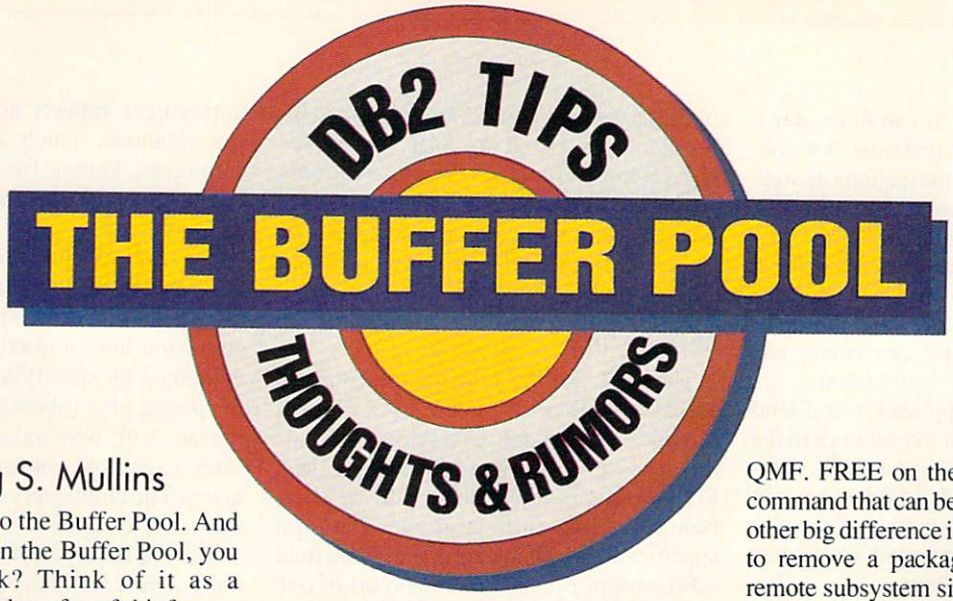


By Craig S. Mullins

Welcome to the Buffer Pool. And what is in the Buffer Pool, you may ask? Think of it as a hodgepodge of useful information, floating around out there in memory, just like DB2's bufferpool. The Buffer Pool will appear regularly in *RELATIONAL DATABASE JOURNAL* and will contain all sorts of interesting information about DB2. Some of it will be in the form of tips, tricks and techniques. Other items will just be rumors — but fun rumors about the future direction of DB2. Still other topics will be random thoughts about DB2, sort of like what you would expect from Andy Rooney if he were a DBA!

### New Release

Well, IBM has officially announced Version 3.1 and it will include increased partition independence, additional and faster bufferpools and DRDA DUOW support. A redesigned index manager was rumored to be included in this new release but, sadly, it is not to be. Although a reliable source indicates that the index manager code is complete, neither the index manager nor row-level locking will be available until Version 3.2. And that is probably about three years down the road.

### DB2 Cycles

According to IBM, there have been two cycles for DB2 development so far and a third cycle is just beginning. Each of these cycles has corresponded to a different version of DB2. Cycle one, lasting from 1984 to 1987, was for Version 1 of DB2 and concentrated on making relational database technology viable in a production environment. The second cycle, corresponding to DB2 Version 2, concentrated mainly on performance and on the advent of client/server and distributed database technology. This cycle lasted from 1988 through 1992. In late 1992 the third cycle began. The focus of this development cycle will be multifaceted with the major impetus on enhanced performance and availability through increased parallelism and technology exploitation. Hmmm, will DB2 utilize hiperspaces soon?

### Why Did They Do That?

Has anyone noticed that the SQL statement to specify a collection-id for an application program is SET CURRENT PACKAGESET? Why not just SET CURRENT COLLECTION? Did a different team develop packages than developed the SET CURRENT PACKAGESET statement? Packages are confusing enough without having to learn two new terms for the same thing!

### A Catalog With A View

Most of you have probably migrated to Version 2.3 by now. Did everything go smoothly? Do you have any views on your DB2 catalog tables? If so, make sure your views are still valid. IBM's catalog migration process does not drop views on the DB2 catalog tables. It can change table definitions though. A good example is SYSIBM.SYSRESAUTH. Under DB2 2.2 the NAME column was eight bytes long. However, under DB2 2.3 it was lengthened to 18 characters to support collections. Any views containing the NAME column of SYSIBM.SYSRESAUTH that were not dropped and re-created will only be able to access an eight-byte NAME. Rule of thumb: When migrating to a new release of DB2, always drop and re-create your DB2 catalog views.

### FREE vs. DROP

What is the difference between freeing packages and dropping packages? Almost nothing. Both will remove the specified package from the DB2 subsystem. Why two different ways? Well, DROP is an SQL statement that can be placed into programs as embedded SQL or issued from SPUFI and QMF. FREE on the other hand is a DSN command that can be issued from DB2I. The other big difference is that FREE can be used to remove a package from any available remote subsystem simply by specifying the appropriate location. DROP removes packages from the current server only.

### More Changes

You may not want to become too dependent on DB2 2.3's catalog visibility feature. The latest rumor is it will not be there in Version 3.1. Not that it will be missed by most DB2 developers. In fact, following is a line that provides the best review of this *wonderful* feature: "It may be slow, but at least it's hard to use."

Does everybody know what a matching index scan is? How about a nonmatching index scan? Well, it looks like IBM is changing terminology again. Remember a few years ago when IBM changed sargable and nonsargable to Stage 1 and Stage 2? Now you can prepare for absolute positioning and relative positioning. These terms are creeping into IBM's DB2 presentations and it would not be surprising to see them in the manuals soon. By the way, absolute positioning is the same as a matching index scan and relative positioning is a nonmatching index scan.

### Explanations

Much confusion surrounds index lookaside. Just how does it work? Index lookaside does not look at the next leaf page, reports a reputable source. First it checks the current leaf page and then it checks the nonleaf page immediately above that leaf page. If you need more information about index lookaside, consult Lockwood Lyon's excellent synopsis of features in his article "DB2 2.3: What Else Is New?" (*DB2 JOURNAL*, August 1992).

Did you know that a latch takes about one-third the number of instructions as a lock. And latches do not need the IRLM. No wonder they are more efficient than a lock.

Look for DB2 3.1 to support data com-

*See The Buffer Pool on page 47*

*Continued from page 29*

ORACLE 7 has many performance enhancements, administration enhancements and SQL*DBA enhancements and changes. By providing shared SQL areas, sites have seen significant performance boosts. The parsed SQL statement, including its execution plan, can be reused by applications utilizing the same command. With optimal caching parameters, multiuser application systems will achieve performance enhancements. Cost-based query optimization in addition to rules-based query optimization can improve the performance of some applications.

## Conclusion

ORACLE provides features familiar to DB2 (and other RDBMS) users. Some of these familiarities include tables, views, indexes, data type support and SQL semantics. Many differences exist, as explained in the body of this article.

ORACLE 7 provides the functionality demanded by government and industry, especially with the introduction of two-phase commit and snapshots (sync replication, read-only at the remote node). ORACLE's trusted server provides extra security features, in addition to the more than 60 available privileges, for environments requiring increased measures of security. ORACLE 7 incorporates many of the wish-list items from the ORACLE user group and customer sites, which reflects the fact that the Oracle Corporation has been listening and has been effective in transforming many ideas into a robust, world-class product.

*ABOUT THE AUTHOR*

*Lawrence Proman is an independent consultant and a JAD, RAD and ORACLE instructor for the Systems Group (Glastonbury, CT). He has 17 years data processing experience as a systems programmer, programming manager, project manager, systems architect and consultant. 3717 Custis Rd., Richmond, VA 23225, (804) 272-1938.*



*Continued from page 30*

pression in hardware much like DB2 2.3 supports sort in hardware.

### *Trivia Question*

Before closing this first installment of the Buffer Pool, here is a trivia question.

The predecessor of DB2 was IBM's RDBMS research project named System R. What does the *R* stand for? Hint: It may not be what you think! Replies can be sent to the author via CompuServe (70410,237) or Prodigy (WHNX44A).

*ABOUT THE AUTHOR*

*Craig S. Mullins is a member of Platinum Technology's technical advisory group and has had extensive experience in all facets of database systems development. He is the author of* DB2 Developer's Guide (ISBN 0-672-30191-1) *published by Sams Publishing (Carmel, IN). Platinum Technology, 1815 S. Meyers Rd., Oakbrook Terrace, IL 60181, (708) 620-5000.*

# BOOK REVIEW

By Jonathan Sayles

# DB2 Developer's Guide

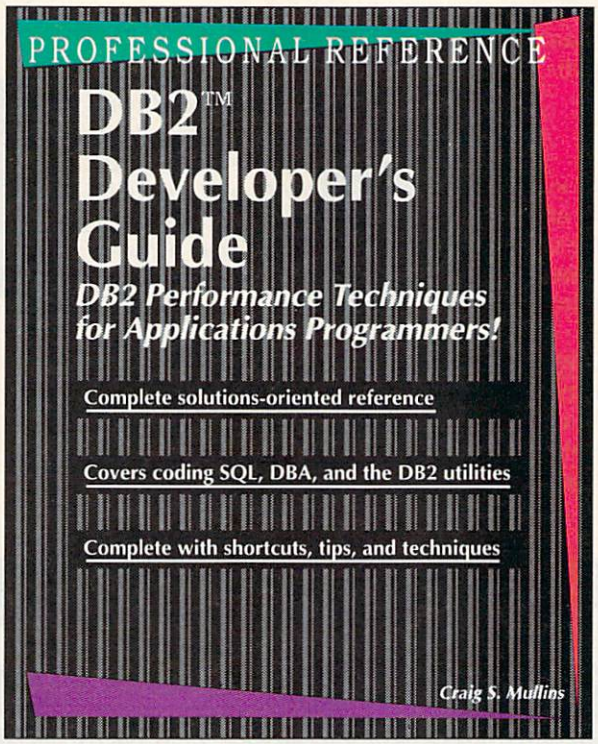


This is an excellent book written by Craig Mullins, one of America's premier DB2 DBAs and internals gurus. It is chock full of DB2 technical information, design and tuning approaches and database administration guidelines.

"DB2 Developer's Guide" presents literally everything programmers and DBAs need to know about advanced DB2. It includes product concepts, DML, DDL and table design guidelines, DCL and security, application programming, dynamic SQL, program preparation and alternative DB2 development tools. DB2 internals are included with in-depth coverage of the optimizer, the system catalog tables, DB2's locking mechanism, DB2 performance monitoring, EXPLAIN and access paths. Other areas covered are tuning the DB2 environment (including MVS, IMS and CICS), the DB2 utilities, DB2 commands, DB2 standards and, believe it or not, much more.

How much more can any single book contain about DB2, you may well ask? Perhaps this will put the mountains of information in this encyclopedic reference into perspective: out of 1123 pages, the previously mentioned topic list ends on page 828. This is not a typo, there are actually 1123 pages in this book. Maybe that fact will make you feel better about paying $59.95 for it. Yes, this could be the longest DP book ever, but the author, in a sort of literary tour de force, does not let up in intensity or quality from start to finish.

In an organized and coherent way, Mullins seems to have dumped his entire DB2 life experience into "DB2 Developer's Guide." The level of detail and technical depth discussed in virtually every topic is explicit and fastidiously accurate. Even more impressive is his style. Mullins manages to describe highly technical points without losing the reader. This is probably based on communication skills, fine-tuned through repeated DB2 subject teaching experiences (where instructors generally have to be able to get their points across in "real time"). Pictures and illustrations, facts and detailed lists, guidelines and rules of thumb—all abound. Perhaps more important than the facts and information in the book are Mullins' opinions, many of which cut against the grain of IBM reference manuals as well as conventional DB2 wisdom (e.g., see the appendix on why not to create base table views). These opinions, it seems, are drawn straight from many real-world experiences and are worth pondering.

**Perhaps more important than facts are Mullins' opinions, many of which cut against the grain of IBM reference manuals.**

Two points concerning this book are worth consideration for a second edition. First, the title suggests an applications orientation. "DB2 Developer's Guide" does include a reasonable amount of applications-oriented material. However, it contains as much, or more for DBAs, and is one of the few books that touches on systems programmer DB2 topics. At least half of the book dwells on DBA subjects such as design, utilities, commands, system-level tuning, etc. Not that applications programmers wouldn't benefit from an understanding of DBA topics, it is just that DBAs would benefit even more. The publisher (Sams Publishing, Carmel, IN) should consider retitling this work to something more inclusive like "DB2 Designer, Developer and Administrator Guide."

Second, (not to make an already lengthy book even longer), there was a lack of anything substantial on distributed database. DDF, DDCS/2 and elements of distributed design, coding and administration would be a welcome addition to the second edition. If Mullins treats distributed DB2 and associated topics, such as client/server, with the same expertise and bravado as he treated the current edition's topics, we would all be the better for it.