

Another issue, another collection of unique DB2 ideas. But first, the answer to last issue's trivia question: The *R* in System R stands for Ruffus. It seems that Ruffus was the name of one of the developer's dogs. Which developer? Sorry, but IBM will not reveal that. It is still classified as proprietary information.

Hear Ye, Hear Ye

If you have not migrated your DB2 subsystems to Version 2.3 yet, you should do so as soon as possible. The only version of DB2 that will be serviced after September 30, 1993 is DB2 2.3. Support for all prior versions and releases of DB2 will be dropped by IBM after this date and migration to Version 3 can be accomplished only from Version 2.3.

Read Efficiency

Most DBAs know the fabled formula for determining read efficiency within a DB2 subsystem. Read efficiency can be defined as the average number of pages that DB2 can request without incurring additional I/O and is calculated by using this formula:

$$\frac{\text{(Total GETPAGES)}}{\text{(Seq Prefetch I/O + Sync Read I/O)}}$$

These numbers are easily obtainable from most performance monitors. Usually, the higher the read efficiency, the better performance will be. However, it is entirely possible that internal DB2 changes could significantly impact this number. For example, consider the DB2 2.3 feature, index lookaside. When index lookaside is utilized, the number of GETPAGES performed by DB2 can be dramatically reduced. So, the read efficiency may go down, but performance will be better. This is because the same number of I/Os are done, but fewer GETPAGES are required due to internal DB2 changes. New versions, releases and even intermediate PTFs possibly could reduce DB2 GETPAGES causing this type of effect. The moral of the story is: *Do not blindly rely on simple calculations to estimate DB2 performance.*

All The Young DB2s

Carry the news — IBM is planning to provide a version of DB2 for every IBM

platform, and even some non-IBM platforms, too. By now you may have heard about the new OS/2 DBMS called DB2/2. This is actually the next version of OS/2 EE Database Manager. It was renamed because DB2 carries with it a high recognition factor and the aura of a successful product. IBM has also announced DB2/6000 for its RISC-based, RS/6000 workstation. And even more surprising, rumors are that a release of DB2/NT for Microsoft Windows NT will be circulating before year-end.

Versions And Releases And Mods

Oh, my! The latest and greatest DB2 is officially called DB2 Version 3 by IBM. It is not called Version 3.1 because of the somewhat cryptic naming conventions used by IBM. You see, each IBM software product has a version and possibly a release and modification level. Versions start at 1 and increase sequentially. A new version does not have a release associated with it until the second release. So, releases start at 2, but the previous release will then be referred to as Release 1. Each release can have a modification level (mod) that starts at 0. Furthermore, releases of 1 and mods of 0 can be dropped from the name. Clear as mud, huh? For simplicity's sake, though, DB2 Version 3 and DB2 3.1 are really the same.

It's About TIME

If you have been trying to determine the last time you executed RUNSTATS, help is on the way with DB2 3.1. Each DB2 catalog table updated by RUNSTATS will have a timestamp column, STATSTIME, indicating the last date and time that RUNSTATS updated that particular table.

And speaking of the DB2 catalog and DB2 3.1, the column distribution statistics will finally be removed from the SYSFIELDS table and placed into a table of their own. Wow, does this mean the DB2 catalog is finally in first normal form? Well, almost. You see, there still remains the matter of all those repeating groups in SYSCOPY and the security tables. But, they are not that much trouble when compared to the SYSFIELDS fiasco.

Security Anomaly

Do you know why a user with the BIND privilege on a plan can free that plan but a user with the BIND privilege on a package cannot free that package? To free a package, the user must meet one of the following conditions: 1) be the owner of the package or 2) have SYSADM or SYSCTRL authority or 3) have BIND-AGENT privilege granted by the package owner. Simple? NOT!

Records Or Rows?

Is a RID a record ID or a row ID? Surprise, surprise! It's a record ID. You see, a DB2 record is the combination of the record prefix and the row. Each record prefix is six bytes long. And RIDs point to the record, not the row. Therefore, a RID is a record ID. But don't let this change the way you think. Those little buggers returned by your SELECT statements are still rows!

Read-Only Relief

Did you know that DB2 2.3 tablespaces started for read-only access will only take tablespace locks? This is true even if the LOCKSIZE of the tablespace is PAGE or ANY. So, if you have any applications that access true read-only tables, consider starting those tablespaces ACCESS(RO).

Trivia Question

Most of you are familiar with the EXPLAIN process and the PLAN_TABLE that goes along with it. Recall, if you will, how a tablespace scan is indicated in the PLAN_TABLE. The value of the ACCESS-TYPE column is set to the character R. Why? No R appears in the words tablespace scan. The question is: What does this *R* stand for? Hint: it's not Ruffus.

Are you sitting on a piece of crucial DB2 information? Does it hurt? Well, get off that thing and share it with the world! How? Send it to The Buffer Pool in care of *Relational Database Journal* and make sure it receives the attention it deserves. Or, contact the author on Prodigy (WHNX44A) or Compu-Serve (70410,237). ☺