



Change Control for DB2 Access Paths

BY CRAIG S. MULLINS

An important aspect of efficient DB2 operation is the ongoing analysis and management of DB2 access path changes. Whether you are implementing changes into your DB2 applications, upgrading to a new version of DB2, or simply trying to achieve optimum performance for existing application plan and packages, an exhaustive and thorough BIND management process is a necessity.

However, many organizations are not doing everything possible to keep access paths up-to-date with the current state of their data. There are several reasons why the acknowledged “best practice” of REOG/RUNSTATS/REBIND is not followed religiously. In this month’s column, I will examine some of the issues involved in managing access path changes. Furthermore, we’ll look at some methods for introducing a better change control discipline for the DB2 access path generation process.

SOME HISTORY

On the mainframe, change has traditionally been strictly controlled. But one exception has been DB2 access paths.

Think about it. In a mainframe shop, everything we do is tightly controlled. If we make even a minor change to an application program, that program is thoroughly tested before it ever reaches a production environment. The program progresses through unit testing, QA testing, volume testing, and so on. As developers, we do a good job of testing a change to minimize the risk that the change might have unintended consequences. We do the same type of due diligence with most other changes in the mainframe world. Database changes are planned and thoroughly tested. System software, e.g., CICS, WebSphere, etc., including subsystem and DB2 changes, are all subject to strict change control procedures. This is done to minimize disruption to the production work being conducted by our business folks.

But there is one exception to this tight change control environment: Binds and Rebinds are typically done in the production environments without the benefit of oversight or prior testing. This lack of change control results in unpredictable performance impacts. In most shops, programs are moved to production and bound there. Indeed, we are at the mercy of the DB2 optimizer, which generates access paths on the fly when we Bind or Rebind

our programs. Any issues with inefficient access paths are then dealt with in a reactive mode. That is, problems are addressed after the fact.

One of the biggest reasons for not implementing strict change control processes for access paths is the lack of built-in methods for ensuring access path change control discipline. Let’s face it, manually evaluating thousands of packages and tens of thousands of SQL statements can be quite impractical. But there are things that can be done to help alleviate this problem. This article will address some of those things.

BIND PARAMETERS

There are many parameters and values that must be chosen from and specified when you bind a DB2 application program. The vast array of options at our disposal can render the whole process extremely confusing – especially if you don’t bind on a daily basis. And even if you do, some of the options still might be confusing if you rarely have to change them. You know what I’m talking about, parameters like ACQUIRE, RELEASE, VALIDATE, and DEGREE.

It is not the intent of this article to delve into the myriad bind options and give you advice on which to use when. There are many articles and books, as well as the IBM DB2 manuals that you can use to guide you along that path. Suffice it to say, that there are some standard parameters and values that should be chosen “most of the time” in certain situations. As such, a wise DBA group will set up canned routines for the programmers to use for compiling and binding their applications. Choices such as: “CICS transaction,” “DB2 batch,” or “analytical query” can be presented to the developer and then, based on which of the various types of programs and environments that are available, the canned script can choose the proper bind options. Doing so can greatly diminish the problems that can be encountered when the “wrong” parameters or values are chosen at bind time.

This same process can be put in place for production binding to ensure that the appropriate parameters and values are chosen. This is especially useful when the binds are not done by a DBA, but are automated in production or done by a less-experienced change control clerk.

Of course, there should always be a method for overriding the “standard” values for special situations, although these overrides should not be available to anyone other than a well-trained individual (DBA or otherwise).

I want to make one small exception here regarding advice on bind parameters, and that is the EXPLAIN parameter. In production, always bind your plans and packages specifying EXPLAIN YES. Failing to do so means that access paths will be generated, but you will not know what they are. This is akin to blinding yourself to what DB2 is doing and is not advisable.

APPROACHES TO ACCESS PATH MANAGEMENT

OK, so we know that Bind and Rebind are important components in assuring optimal application performance. It is the bind process that determines exactly how your DB2 data is accessed in your application programs. As such, it is critically important that you develop an appropriate strategy for when and how to rebind your programs.

There are several common approaches taken by DB2 users. By far, the best approach is to rebind your applications over time as the data changes. This approach involves some form of regular maintenance that keeps DB2 statistics up to date and formulates new access paths as data volumes and patterns change. There will be more on this later on in this article.

Other approaches include binding only when a new version of DB2 is installed, or perhaps more ambitious, whenever new PTFs are applied to DB2. Another approach is to rebind automatically after a regular period of time, whether it is days, weeks, months, or whatever period of time you deem significant. This approach can work if the period of time is wisely chosen based on the application data – but it still can pose significant administrative issues.

The final approach is from the “if it ain’t broke, don’t fix it” school of thought. This approach is the worst of the several approaches discussed here. The biggest problem with this approach is that you are penalizing every program in your subsystem for fear that a program or two may have a few degraded access paths. This results in potentially many programs having sub-optimal performance because the optimizer never gets a chance to create better access paths as the data changes.

Of course, the possibility of degraded performance is real – and that is why this approach has been adopted at some sites. The problem is being able to find which statements may be worse. The ideal situation would be to be able to review the access path changes before hand to determine if they are better or worse. But DB2 itself does not provide any systematic method of administering access paths that way. There are third-party tools that can help you achieve this, however.

Let’s go back to the best approach again, and that is to perform regular rebinds as your data changes. This involves what has become known as the “Three Rs.” This means regularly reorganizing the data to ensure that it is optimally structured. That is followed by RUNSTATS to be sure that the reorganized state of the data is reflected in the DB2 Catalog. Finally, we follow that up with rebinds of the application programs that access the data structures that have been reorganized and RUNSTATEd (if you’ll allow me to turn that into a verb).

At any rate, your goal should be to keep your access paths up-to-date with the current state of your data. Failing to do this means that DB2 is accessing data based upon false assumptions. DB2 is

unlikely to make the same access path choice as your data grows – and as patterns within the data change.

By rebinding, you can generally improve the overall performance of your applications because the access paths will be designed better, based on an accurate view of the data. Additionally, as DB2 changes are made (via new releases or PTFs) optimizer improvements, and new access techniques can be incorporated into the access paths. That is, if you never rebind, not only are you forgoing better access paths due to data changes but you are also forgoing better access paths due to changes to DB2 itself.

Of course, adopting the Three R’s approach can pose additional questions. For example, when should you reorganize? In order to properly determine when a REORG is needed, you’ll have to look at statistics. This means looking at either RUNSTATS or Real-Time Statistics (RTS). So, perhaps it should be at least Four R’s – in other words:

1. RUNSTATS or RTS,
2. REORG,
3. RUNSTATS, and
4. REBIND.

Now it is true that some folks don’t rely on statistics to schedule a REORG. Instead, they just build the JCL to REORG their database objects when they create the object. So they create a table space then build the REORG job and schedule it to run monthly, or quarterly, or on some regular basis. This is better than no REORG at all, but it is probably not the best approach because you are most likely either reorganizing too soon (in which case, you waste the CPU cycles to do the REORG) or you are reorganizing too late (in which case, performance is suffering for a period of time before the REORG runs). Better to base your REORGs off of statistics and thresholds using either RUNSTATS or RTS.

Statistics are the fuel that makes the optimizer function properly. Without accurate statistics, there is little hope that the optimizer will formulate the best access path to retrieve your data. If the optimizer doesn’t have accurate information on the size, organization, and particulars of your data then it will be creating access paths based on either default or inaccurate statistics. Incorrect statistics will probably cause bad choices to be made – such as choosing a merge-scan join when a nested loop join would be better, or failure to invoke sequential prefetch, or using the wrong index – or no index at all. And the problem of inaccurate statistics is pervasive. There are shops out there that never, or rarely, run RUNSTATS to gather up-to-date statistics. Make sure yours is not one of those shops!

When should you run RUNSTATS? One answer is: “As frequently as possible based on how often your data changes.” This means that you will need to know a thing or two about your data growth patterns. To properly determine a schedule for statistics you need to know things about your data: what is its make-up, how is it used, how fast does it grow, and how often does it change? These patterns will differ for every table space in your system.

Next we need to decide when to rebind? The best answer for this is when statistics have changed significantly enough to change access paths. When we know that data has significantly changed, it makes sense to rebind after the RUNSTATS completes. But the trick is determining exactly when we have a “significant” change in our data. Without an automated method of comparing and contrasting statistics (or even better yet, access paths) coming up with an answer in a manual way can be time-consuming and error-prone – especially when we get into the thousands of programs.

And we always have to be alert for a rogue access path – that is, when the optimizer formulates a new access path that performs worse than the previous access path. This can happen for a variety of reasons. Of course, number one is that the optimizer, though good, is not perfect. Mistakes can happen. Other factors can cause degraded access paths, too. The access paths for volatile tables depend on when you run the RUNSTATS. Volatile tables are those that start out empty, get rows added to them during processing, and are emptied out at the end of the day. And, of course, if the catalog or statistics are not accurate we can get problems, too.

So adopting the Three, err, I mean, Four R's approach implies that you will have to develop a methodology for reviewing your access paths and taking care of any "potential" problem access paths. This can be a difficult mountain to climb.

Indeed, the Four R's probably needs to become the Five R's because we need to review the access paths after rebinding to make sure that there are no rogue access paths. So, we start off with a RUNSTATS (or use RTS) to determine when to REORG. After reorganizing we should run RUNSTATS again, followed by a REBIND. Then, we need that fifth R – which is to "review" the access paths generated by the REBIND. As we mentioned, the optimizer can make mistakes. And, of course, so can you. Users don't call you when performance is better (or the same). But if performance gets worse, you can bet on getting a call from irate users.

So we need to put in place best practices whereby we test bind results to compare the before and after impact of the optimizer's choices.

THE PLAN TABLES

A lot of information is contained in the PLAN_TABLE. After the optimizer creates the access paths and populates the PLAN_TABLE with data representing those access paths, we need to examine the results to determine if everything is OK.

Many questions can be answered by analyzing the results of EXPLAIN – questions like:

- If we are joining, what type of join is used (NLJ, MS, Hybrid)?
- Was an index used, and if so how many columns matched?
- Are we doing a scan, and if so, what type of scan (full or page range)?
- Is prefetch being used, and if so what type (sequential, list)?
- Was a hint used?
- Was parallelism used, and if so, what degree and type (I/O, CPU, Sysplex)?
- Was a sort required, and if so, why (Join, Unique, Group By, Order By)?
- What type of locking is required?

And that just covers the main PLAN_TABLE. The EXPLAIN option also populates two optional tables, if they exist:

- DSN_STATEMENT_TABLE, which contains DB2's estimate of the processing cost for an SQL statement, and
- DSN_FUNCTION_TABLE, which contains information about function resolution.

Of course, for any of this information to be returned, you have to have bound specifying EXPLAIN(YES). Any change to any of these items between rebinds means a change in access path – which can be positive, or a potential problem. Over time, performance

analysts can determine which changes are good and which might be problematic – but it takes experience (and perhaps some luck) to do this correctly. Using a tool that automates the process can also make the task much easier and more accurate.

So, how do you determine what access paths have changed? Sometimes the program has changed, too – which can make it challenging to find the exact SQL statements to compare. When just the access paths change, it will be easier to compare them and spot the changes, but there is still a wealth of data that needs to be analyzed to do this justice.

And when you are talking about thousands of programs being rebound, do you really have the time to review every access path to make sure it is fine? This question alone causes many folks to go back to the "Let it Ride" mentality – which is too bad, because it is an inferior approach, especially when there are products that can help.

VERSION MIGRATION ISSUES

Let's switch gears and talk about an impending event that many of us are still facing, namely migrating from DB2 V7 to V8. First of all, let's be clear, you do not have to rebind all of your packages and plans when you move to V8. But it is a really good idea to do so, and most of you will probably rebind most, if not all, of your programs when you get to V8. Why?

First of all, there are optimizer and performance improvements that you won't get without a rebinding. And there will be degraded program performance that will occur when you get to V8 that rebinding can fix. And for some of you, there will even be REBINDs that you just will not be able to avoid. Let's examine each of these issues briefly.

First of all, what is the "degraded performance" issue? The problem occurs when DB2 turns off fast column processing. DB2 V3 introduced a function called an SPROC. An SPROC, or SELECT procedure, enables fast column processing. Essentially, this enhancement examines SELECT statements that are executed repeatedly and builds an internal procedure that moves all the columns in one move rather than one column at a time. You have no external control over when or if DB2 uses them. And the more columns that are specified on a SELECT, the greater the performance gain could be.

How does this all tie into Version 8? If a plan or package is using an SPROC in V7, the SPROC is using 31-bit code. When you attempt to run that same plan or package in V8 without rebinding it first, it needs to be in 64 bit. It isn't, so DB2 disables the procedure. The only way you can re-enable the SELECT procedure is by rebinding the program. Until you do that rebinding, and if the plan or package uses an SPROC, your application's performance will be degraded. Do the rebinding, and you should see a performance improvement. Along those lines, the IBM redbook titled "DB2 UDB for z/OS Version 8 Performance Topics" specifically warns of this problem, cites the potential for CPU increases of up to 10 percent and recommends global rebinds.

And what about those rebinds that cannot be avoided. Well, DB2 V8 will autobind any plans and packages that were bound prior to DB2 Version 2 Release 3. So you might experience an execution delay the first time such plans are loaded unless you rebinding them yourself. And DB2 might change the access path due to the autobind, potentially resulting in a more efficient access path – or a more inefficient access path.

Such actions might become more common in future DB2 ver-

sions. In several conference presentations, folks at IBM have suggested that in the future DB2 may autobind any plan or package that was last bound on an “out of service” version of DB2. What might that mean for DB2 V9? Right now, only V7 and V8 are in service, so think about that when you are considering your rebind approach.

And there are still more reasons to rebind when moving to V8. DB2 V8 in NFM uses a different format for its DBDs, packages and plans. Before it can use a DBD, plan or package from an older DB2, it must first be expanded to the new Version 8 format. This causes more overhead. What should you do? Here is the advice right out of the afore-mentioned redbook:

After you have entered new-function mode, we recommend that you plan to rebind all of your plans and packages. DB2 will then store the plans and packages in the DB2 catalog in the new format. DB2 will no longer need to expand the plans/packages each time it needs to use them.

SUMMARY

Forward-thinking organizations should adopt a liberal Bind/Rebind process to ensure optimal access paths based on up to date statistics. Keeping abreast of data changes and making sure that your programs are optimized for the current state of the data is the best approach. This means regular executions of RUNSTATS, REORG, and REBIND. If you are worried about rogue access paths, consider investing in a third-party tool that can assist with access path changes management issues.

Failing to keep your access paths aligned with your data is a sure recipe for declining DB2 application performance.

ABOUT THE AUTHOR

Craig S. Mullins is president and principal consultant with Mullins Consulting, Inc. He is an IBM Gold Consultant, the author of two books, DB2 Developer’s Guide, 5th ed. and Database Administration: Practices and Procedures, and can be reached via his Web site at www.CraigSMullins.com.

Index to Advertisers

| | |
|-----------------------|--------------------|
| IDUG Conference | 32 |
| HLS Technologies | 3 |
| Relational Architects | Inside Back Cover |
| Responsive Systems | Back Cover |
| BMC Software, Inc. | Inside Front Cover |