

WHITE PAPER

Best Practices for Optimizing Db2 Performance

A Guide for DBA Managers

by Craig S. Mullins



datAvail

BI/Analytics • Applications • Databases

Contents

Introduction.....	1
Technology Trends Drive DBA Management	1
Db2 Performance: At Fifty Thousand Feet	2
Db2 Application Performance	3
Db2 Database Objects and Performance.....	5
Clustering	6
Maintaining Your Space	7
Db2 Subsystem Performance	8
The Complexity of Db2 Performance	12



Introduction

Db2 performance tuning and optimization is a complex issue comprising multiple sub-disciplines and levels of expertise. Mastering all of the nuances can take an entire career. Deploying standard best practices can minimize the effort to achieve efficient Db2 applications and databases.

This white paper outlines the most important aspects and ingredients of successful Db2 for z/OS performance management. It offers multiple guidelines and tips for improving performance within the three major performance tuning categories required of every Db2 implementation: the application, the database and the system.

Along the way, we will look at some of the most common problems that DBAs encounter, and offer tips and guidelines for assuring effective performance of your Db2 databases and applications.

Technology Trends Drive DBA Management

In today's modern enterprise, information is power; and database systems are the predominant technology for storing, managing, and accessing the data that businesses use to glean information.

Db2 is one of the leading database systems on the market. It's used by 100% of Fortune 100 firms and 80% of the Fortune 500. Therefore, it is an important component that drives the most significant businesses in the world.

Dynamic business requirements force constant change and overwhelming complexity in the typical IT environment. Regulatory compliance issues, time-to-market pressures, and industry consolidations cause turbulence as IT struggles to keep systems in sync with business demands. But managing systems becomes even more troublesome as technology trends evolve. Examples of these trends include business intelligence, e-commerce, and rapid software versioning. And, you can expect data growth will continue unabated.

At the same time, the mainframe workforce is aging and retiring. As the number of mainframe professionals with years of experience decrease, they must be replaced with younger technicians – and these technicians require a different interface than the green-screen, ISPF interface that is still the



most common mainframe interface around. According to the Mainframe Report from Arcati, Ltd.: "When a (mainframe) person leaves, it is not simple technical knowledge they take with them but rather twenty years' intimate and detailed understanding of core business applications."¹

These trends drive home the need for DBAs to exert constant vigilance over the performance and availability of their Db2 environment. Any performance slip can impact the business resulting in lost revenue or damaged customer relationships. Furthermore, it is important to be able to achieve this performance management capability with a modern, easy-to-use interface. Enterprises cannot reap the full range of benefits from their database systems without modern, effective performance monitoring and management capabilities.

To improve Db2 performance management, you can use solutions and services that can simplify and automate the performance management process in a cost effective manner. There are many powerful, Db2 performance monitoring tools available, but many are difficult to use. As DBA mainframe skills gap continues to widen, more flexible tools and services can make analyzing Db2 performance easier.

"When a (mainframe) person leaves, it is not simple technical knowledge they take with them but rather twenty years' intimate and detailed understanding of core business applications."

Db2 Performance: at Fifty Thousand Feet

To help meet these challenges, DBAs must possess the necessary tools to automate the maintenance and performance management of their Db2 environment. But what do we really mean by Db2 performance?

Let's start by thinking about this question at a high level – a 50,000 foot level, if you will. Every Db2 application requires three components to operate: the application, the database and the system itself. The application is the host language code and SQL that provides each program's functionality. The database refers to the database objects that house the application data. And the system refers to the Db2 subsystem installation and its interfaces to other system software (e.g. z/OS or CICS). To effectively deliver Db2 performance, DBAs must be able to monitor and tune each of these components.

Let's take a look at some of the most common problems that DBAs encounter when managing each of these components.

¹ The Arcati Mainframe Yearbook 2005, "The Mainframe: Forty Years On," page 7

Db2 Application Performance

The first component of Db2 performance is the application itself. As much as 80% of all relational database performance problems can be tracked back to inefficient application code. The application code consists of two parts: the SQL code and the host language code in which the SQL is embedded.

Although SQL is relatively easy to learn (at the most basic level), it is actually quite difficult to master. There are, of course, some general rules of thumb you can apply to help you to build efficient SQL statements. Keep in mind though, that a rule of thumb is helpful as a general guideline, but shouldn't be applied in every situation.

The first major rule of thumb for SQL programming is to build the SQL to do the work instead of the application program. This means that you should code the proper WHERE clauses in your SQL instead of waiting to filter data in your host program. Additionally, do not treat Db2 tables like master files. Code SQL joins instead of multiple cursors where a "master" cursor is read, that then drives the other cursors. SQL is a set processing language and it should be coded that way.

One common problem is retrieving more data than is required. Your SQL statements should retrieve only the columns required, never more. Each column that has to be accessed and returned to you adds overhead that can be avoided if the column is unnecessary. Sometimes programmers try to use the SELECT * shortcut – and though that may be fine for quick and dirty testing, it should never be allowed in production programs.

It is also important to understand the different types of predicates and their impact on performance.

A predicate that can be satisfied by the Data Manager portion of Db2 is referred to as Stage 1. A Stage 2 predicate has to be passed from the Data Manager to the Relational Data System to be satisfied. The Data Manager component of Db2 is at a level closer to the data than the Relational Data System. The earlier in the process that Db2 can evaluate the predicate, the more efficient processing will be. So, Stage 1 predicates are more efficient than Stage 2.

Additionally, a query that can use an index has more access path options, so it can be more efficient than a query that cannot use an index. The Db2 optimizer can use an index or indexes in a variety of ways to speed the retrieval of data from Db2 tables.

For this reason, try to use indexable predicates rather than those that are not.

At this point you are probably asking "which predicates are Stage 1 and which are indexable?" Where predicates are processed can change from version to version of Db2. This information is documented in the IBM Db2 Managing Performance manual,² which can be downloaded from the IBM web site free-of-charge.

Another technique for improving performance is to create indexes to support your ORDER BY and GROUP BY clauses. If an index is available on the columns you specify to these clauses, Db2 can use the index to avoid invoking a sort. That is likely to improve the performance of your SQL statement.

After the SQL is coded, to ensure proper performance you can review and interpret access path information collected using the EXPLAIN command. The EXPLAIN command can be executed on a single SQL statement or multiple statements in an entire program. EXPLAIN captures information about the access paths chosen by the Db2 optimizer and populates it into a special table called a PLAN_TABLE.³ The PLAN_TABLE contents are encoded and a DBA or analyst must interpret that data in order to determine exactly what Db2 is doing.

"Index tuning can be tricky; changing an index to improve the performance of one query can degrade the performance of other queries."

² For Db2 10 for z/OS the manual number is SC19-2978. Predicate information can be found in Table 66, starting on page 269.

³ Db2 populates more than just the PLAN_TABLE. There actually are more than a dozen explain tables as of Db2 10 for z/OS. A detailed list of these tables can be found in the Db2 Managing Performance manual (SC19-2978) in Appendix B.



To effectively tune SQL you will require additional performance metrics than just what is in the PLAN_TABLE. Based on execution circumstances, Db2 may have to re-evaluate an access path at run time. For example, to avoid using RIDs in the event of a RID pool failure or if there is a lack of storage for a list prefetch operation.

SQL performance tuning requires an ability to interpret the contents of the PLAN_TABLE in conjunction with the SQL code, host language code and information from the Db2 catalog to judge the efficiency of each SQL statement. Analysis of the access path information and additional performance metrics can require a tuner to make SQL modifications. Perhaps you will need to modify predicates to influence indexed access, or revise a statement to use Stage 1 instead of Stage 2 predicates.

An in-depth analysis of your SQL statement can determine that index changes or additional indexes will be required. So, you see, application tuning can lead back to database object tuning. And index tuning can be tricky; changing an index to improve the performance of one query can degrade the performance of other queries.

Build indexes on the predicates of your most heavily-executed and most important queries. Do not limit the number of indexes per table. Create as many as are necessary to improve query performance. But remember to balance that against your data modification needs. An index can improve query performance, but it will degrade the performance of data modification because the index has to be modified when its underlying data is inserted, deleted, or updated.

You might also consider overloading indexes by adding columns to encourage index only access. For example, if you have a query that is frequently run throughout the day and it accesses five columns, four of which are in an index, adding the fifth column to that index allows Db2 to get all of the data from the index alone. Avoiding I/O to the table space can help to improve performance.

Remember that inefficient SQL coding and improper indexing are not the only reasons for poor application performance. An inefficiently designed host language program can cause poor performance, too. Host language code is the code written in a programming language such as Java, COBOL, C, or your programming language of choice. SQL statements are embedded into the host language code and it is possible to have a well-tuned SQL statement embedded into an inefficient application program.

Once again, to effectively manage the performance of your Db2 applications requires flexible, and easy-to-use tools. It is important that you are able to quickly find the programs that are consuming the most resources. You do not want to have to wade through pages and pages of performance reports to find the offending programs; neither do you want to have to navigate through screen after screen of options and confusing menus.

The best option is to use a tool or service that works as a tuning lab for your SQL. You will need the ability to find the offensive SQL and then be able to play “what if” scenarios by modifying the SQL and gauging the performance impact of each change. It is even better if the tool can provide expert tuning advice to guide you as you change your SQL statements.

Db2 Database Objects and Performance

Although many factors contribute to poor database performance, it is particularly important to pay attention to your database objects. The term database objects refers to the table spaces, tables, and indexes that comprise your Db2 databases.

Of course, the first factor for ensuring efficient database objects is practicing proper database design. This means starting from a fully-normalized, logical data model. The data model should be used as a template for building your physical Db2 database, deviating from the model only for performance reasons.

After implementing an effective physical Db2 database, the next step is to ensure that you are capturing statistics appropriately for the database objects. The use of inaccurate or outdated database statistics by the Db2 query optimizer often results in a poor choice of query execution plans and hence unacceptably long query processing times. According to information supplied by IBM's Silicon Valley Lab at a recent IDUG conference, as many as half of all access path "problems" presented to IBM's technical support were caused by inaccurate, outdated, or missing statistics.

Statistics are accumulated by the RUNSTATS utility. You should schedule a regular execution of RUNSTATS for all dynamic database objects. If the data grows or changes on a regular basis, you need to plan and run the RUNSTATS utility accordingly. Only by having up-to-date statistics can the Db2 optimizer have the correct information to generate accurate access paths for your SQL queries.

RUNSTATS also gather additional information about the current state of each object. Of course, Real Time Statistics (RTS) can be used for this purpose, too. Whenever possible, favor using RTS over other Db2 Catalog statistics because the RTS values will be more up-to-date and therefore will more accurately reflect the state of your Db2 objects.

At any rate, organization data can be used by your DBA to determine when an object needs to be reorganized. Table spaces and indexes become disorganized as business users run application programs that insert, update, and delete data in Db2 tables – causing the once-ordered data to become fragmented. Disorganized data can be particularly detrimental to the performance of your Db2 systems. To understand how disorganization impacts performance, let's review the concept of clustering.



⁴Sysplex query parallelism is deprecated in Db2 11 for z/OS.

Clustering

Each Db2 table can have a single clustering sequence defined for it. This is accomplished with a clustering index. When clustering is defined, Db2 attempts to maintain the sequence of the rows in the physical order in the table space.

But as data is inserted and modified and pages fill up, there may not be sufficient space available to insert the data in the proper sequence. So the data is inserted where space exists and clustering begins to break down. As the data becomes unclustered, the performance of sequentially accessing data in clustering order begins to degrade because more I/O operations are required to retrieve the data. Running a reorganization will re-cluster the data, and thereby improve performance.

Clustering is not the only organization detail that can impact performance. You must also keep an eye on other details such as:

- The percentage of each table space containing data from tables that have been dropped. As this percentage increases, performance gets worse.
- The leaf distance for indexes, which indicates the average number of pages between successive index leaf pages. Performance degrades as the leaf distance increases.
- Far-off and near-off pages for indexes, which estimates how many of the rows in the table are inefficiently located. As the number of far- and near-off rows increases, performance gets worse.
- Near and far indirect references for a table space, which indicate the number of rows that have been relocated either near (2 to 15 pages) or far away (16 or more pages) from their original location. This relocation can occur as the result of updates to variable length rows. As the number of near and far indirect references increases, performance gets worse.

All of these factors impact the efficiency of your database objects. Failing to monitor and correct organization problems will cause every application accessing these objects to suffer performance degradation. The wise course of action is to monitor these statistics. Set thresholds for each that cause a REORG to be scheduled when the threshold value is reached. For example, monitor the cluster ratio and when it falls below 90%, schedule a REORG.

Creating the appropriate indexes for your Db2 tables and applications is a discipline that spans database and application tuning. To create appropriate indexes, the DBA must have the SQL that is to be used against the Db2 objects in order to understand the access patterns. Furthermore, you need to understand your system's workload. This information enables you to balance index creation and buffer pool placement so that your setup matches what is being run against it.

Usually though, DBAs tend to create indexes before having knowledge of the SQL that will be running against their tables. Some indexes are required regardless of the SQL being run, for example, to support a primary key or unique constraint, or to enhance the performance of referential constraints. But the vast majority of your indexes should be to support SQL in your application code.



Maintaining Your Space

Of course, there are more severe database performance problems that you can encounter. For example, consider the impact of running out of space to your applications. When a table space runs out of space and has no more room to expand, any attempt to add more data will fail. And a failure is the worst kind of performance problem – an availability problem.

The best approach to managing the performance of your Db2 database objects is to use a tool or service that can help you automate the setup and monitoring of space and organization thresholds. If you fail to take this approach, you are probably either wasting CPU cycles by reorganizing before it is required or you are incurring performance problems because you are not reorganizing until well after it is required. A dynamic, automated toolset can help reorganize at just the right time and thereby optimize your resource usage and your Db2 systems... and to predict outages and problems so you can resolve the issues before they happen.

Db2 Subsystem Performance

Successfully managing Db2 performance also requires monitoring and tuning your Db2 subsystems. In order to deliver consistent system performance, the DBA must have the tools to monitor, manage and optimize the resources used by Db2.

Tasks required for system tuning include the proper allocation and management of memory structures (e.g., buffer pools, EDM pool, etc.), storage management, integration of the DBMS with other system software, proper usage of database logs, and coordination of the operating system resources used by the DBMS. Additionally, the DBA must control the installation, configuration and migration of the DBMS software. If the system isn't performing properly, everything that uses the system will perform poorly. A poorly performing system impacts every database application.

Efficient memory usage is one of the biggest struggles for Db2 performance tuners. Relational database systems love memory and Db2 uses memory for buffer pools, the EDM pool, the RID pool and sort pools. The more efficiently memory is allocated to these structures, the better Db2 will perform.

Db2 provides 80 different buffer pools that can be configured to cache data in memory for Db2 programs to access. Forget about trying to follow a cookie-cutter approach to buffer pool management. There is no simple formula to follow that will result in an optimally buffered system that works for every implementation. Each shop should create and optimize a buffer pool strategy for its own data and application mix.

The trick is to tune your buffer pools so that they match your workload. You also have to ensure that you do not exceed the memory allocation for your operating system and platform, or system paging will occur and performance will suffer.

To configure your buffer pool based on workload, you need to know how your database objects are accessed. There are basically two types of access: sequential and random. An access request is sequential if it starts by reading one record and continues reading the next records in sequence; this may or may not reflect the way the data is actually stored on disk. Random access, on the other hand, is typified by requests that access a single record based on a key.

If you have database objects that are accessed primarily sequentially, you can take advantage of this knowledge by assigning them to a buffer pool that is configured for sequential access. This is accomplished by setting that buffer pool's sequential steal threshold (VPSEQT), the parallel sequential steal threshold (VPPSEQT), and the assisting parallel sequential steal threshold (VPXPSEQT)⁴ to indicate that access is mostly sequential. These thresholds dictate to Db2 how much of the buffer pool should be set for sequential processes.

Consider the VPSEQT parameter, for example. It specifies the percentage of the buffer pool that should favor sequential processing. Take for instance, by setting VPSEQT to 95, Db2 will favor stealing a random page over a sequential page until 95% of the buffer pool is filled with sequential pages. VPPSEQT and VPXPSEQT are specified as percentages of VPSEQT and are used to specify the sequential steal threshold for parallel operations.

"To successfully configure your buffer pools you will need to be able to determine how your database objects are being accessed: sequentially or randomly."

To successfully configure your buffer pools you will need to be able to determine how your database objects are being accessed: sequentially or randomly. Furthermore, you will need to have a method of monitoring the performance of those buffer pools once configured. Each buffer pool has additional thresholds associated with it that provide usage and control information about buffer pool operations.



You will want to monitor the prefetch disabled threshold, especially for buffer pools that are tuned for sequential access. This threshold kicks in when 90% of the buffer pool pages are unavailable. When this threshold is reached, Db2 no longer initiates prefetch operations and will cancel prefetch operations in progress.

Additionally, you will need to watch for the data manager critical threshold. It is reached when 95% of the buffer pool pages are unavailable. At this point Db2 will access pages once for each row that is retrieved or updated in that page. This should be avoided at all times. Basically, when you hit this threshold and you retrieve or update more than one row on the same page, Db2 will invoke a separate page-access operation for each access that will cause severe performance degradation.

Finally, you will need to monitor the immediate write threshold. When 97.5% of the buffer pool pages are unavailable, Db2 writes updated pages as soon as the update is complete. In other words, writes are performed synchronously instead of asynchronously. This causes an increase in the number of write I/Os performed by Db2.

If any of the last three thresholds are reached, there are only two methods of making changes to alleviate the problem: either increase the size of the buffer pool or reallocate the database objects to other buffer pools. Remember, larger buffer pool sizes are not always more efficient.

Db2 is not the only consumer of memory in your system, so take care when allocating buffer pool memory.

A good buffer pool monitoring technique is to monitor the buffer pool hit ratio for each of your buffer pools. The hit ratio can be calculated as follows:

$$\frac{((\text{GETPAGES} - \text{SUM OF ALL PAGES READ}) / \text{GETPAGES}) * 100}{}$$

The “sum of all pages read” must account for synchronous and asynchronous I/O. This ratio tells you the percentage of times that a requested page was available in memory. When a page is available in memory Db2 does not have to read the page from disk and performance improves. The larger the hit ratio number, the better.

Up to this point we have been more worried about reading data, than writing it. But there are thresholds that need to be tuned that control how modified data is written from the buffer pools back to disk. These are the deferred write thresholds (DWQT and VDWQT).

DWQT is set as a percentage of the buffer pool that can be occupied by updated pages and in-use pages. The default value is 50%, but this is usually too high for most sites. When the percentage of unavailable pages exceeds the threshold, asynchronous writes are scheduled for the updated pages. VDWQT is similar to DWQT, but for a single data set; its default is 10% (and that is likely to be too high as well).

By lowering your deferred write thresholds your Db2 subsystem will have regular asynchronous writes of updated pages to disk. Pages that are frequently referenced will remain in the buffer pool even if updates are written out to disk. But never set VDWQT higher than DWQT; doing so would have no effect because the pages for each data set also apply to the overall percentage.

Another consumer of memory in your Db2 subsystem is EDM storage, which is used for caching internal structures used by Db2 programs. EDM storage is composed of five components, each of which is in a separate storage area:

- 
- EDM RDS pool below: A below-the-bar pool, which contains the part of the cursor tables (CTs) and package tables (PTs) that must be below the bar.
 - EDM RDS pool above: An above-the-bar pool that contains the part of the PTs and CTs that can be above the bar.
 - EDM DBD pool: An above-the-bar pool that contains database descriptors (DBDs).
 - EDM statement pool: An above-the-bar pool that contains dynamic cached statements.
 - EDM skeleton pool: An above-the-bar pool that contains skeleton package tables (SKPTs) and skeleton cursor tables (SKCTs).

As a general rule, you should shoot for an 80% hit rate for DBDs and skeletons in EDM storage. Another way of stating this is that only one out of every five times that a DBD/skeleton is required should it need to be loaded into memory.

It is important to monitor EDM usage because an EDM Pool failure can bring Db2 to a screeching halt. If there is no place to cache the structures needed for a new iteration of a program to run, then that program will wait. And so will you.



Memory allocation is not the only component of Db2 subsystem tuning. You will also need to be able to view information about locking to resolve lock timeout and deadlock situations, as well as to be able to configure the locking system parameters (DSNZPARMs) appropriately.

Actually, having a method to quickly and easily view the DSNZPARM settings for your Db2 subsystem is a necessity – especially for those times when you are in the middle of resolving a difficult problem.

Determining the cause of locking problems can be particularly troubling. When wait time is excessive, performance will degrade. A vigilant DBA Manager must be capable of tracing the lock list to determine if it is properly sized, checking the average time that applications are waiting (and why), and figuring out how many locks are allocated for a database.

Keep in mind, too, that this is a high-level paper on Db2 performance management but there are many more components required to achieve proper Db2 system performance. Other system elements requiring attention include allied agent (CICS, TSO, etc.) tuning, monitoring disk I/O, tuning logging, and Parallel Sysplex configuration and management for Db2 data-sharing shops.

Keeping an eye on so many performance monitoring and tuning details can be an overwhelming task without a tool that can gather and visually represent the diagnostic information in an organized fashion.

The Complexity of Db2 Performance

We have discussed a number of the most common Db2 performance-related issues here. After digesting this information, you know that Db2 performance management is quite complex. Mastering all of the nuances covered in this paper is difficult – and remember, these are just the most common problems. Much more needs to be mastered to ensure efficient Db2 performance.

As the dynamics of the workforce change and new releases of Db2 add features and complexity, the problem is exacerbated. The bottom line is that software tools that automate the detection and correction of Db2 performance problems helps simplify Db2 performance management.

Furthermore, database performance tools need to be designed to be easy to use, yet highly functional. As skilled mainframe professionals retire, they will be replaced with younger technicians who have not had the same level of exposure to original mainframe systems. These younger technicians will expect to use tools with graphical interfaces, pull-down menus, and point-and-click functionality. Of what possible use is a complex tool with broad functionality that no one can figure out how to use?

Conclusion

To reach your Db2 performance management objectives you must be able to monitor and tune all three major components of your Db2 environment: your Db2 applications, database objects, and Db2 subsystems.

A wise DBA manager will deploy integrated performance tools and augment them with skilled services to ensure the efficient performance of the Db2 environment. Doing anything less will cause performance problems... and all of the related business problems that go along with them.

For more information on Db2 performance tuning or managed services call Datavail toll-free at (866) 828-7843. Prefer to chat online? We have experts available 24x7x365 to answer your questions on our chat line. Or let one of our experts call at your convenience: just email us with your phone number and a good time to talk, and we'll call you.



Biography

Craig S. Mullins



Craig S. Mullins is a data management strategist, researcher, and consultant working with Datavail and its Db2 practice to expand offerings and client base.

He is President and Principal Consultant at Mullins Consulting, Inc. and the publisher of The Database Site (thedatabasesite.com). He has three decades of experience in all facets of database management and has worked with Db2 on the mainframe since V1.

About Datavail


Datavail is a company of over 1,000 professionals helping clients build and manage applications and data via a world-class tech-enabled delivery platform and software solutions across all leading technologies.


Contact Us

Corporate Headquarters

Datavail Corporation
11800 Ridge Parkway
Suite 125
Broomfield, CO 80021

General Inquiries

 877-634-9222

 303-469-2399

 info@datavail.com

India

Powai Office

Datavail Infotech Pvt. Ltd,
A-902, Supreme Business Park,
Hiranandani Gardens, Powai, Mumbai – 400076, Maharashtra

Bangalore Office

Datavail Infotech Pvt. Ltd
Ground floor, South Wing, Maruthi Chambers, Rupena Agrahara,
Hosur Main Road, Bommanahalli, Bengaluru – 560068

Hyderabad Office

Datavail Corporation
1002 – 1003, Manjeera Trinity Corporate, JNTU-Hitech City Road,
K P H B Phase 3, Kukatpally, Hyderabad, Telangana-500072

Sri Lanka

Datavail Corporation
No 341/5, Level 4, M&M Center, Kotte Road, Rajagiriya 10100, Sri Lanka

datAvail

BI/Analytics • Applications • Databases

www.datavail.com | 877.634.9222