SQL Performance and Tuning

DB2 Relational Database June 2002 Penny Bowman and Rick McClendon



- The DB2 Optimizer
- SQL Coding Strategies and Guidelines
- DB2 Catalog
- Filter Factors for Predicates
- Runstats and Reorg Utilities
- DB2 Explain
- DB2 Insight

The DB2 Optimizer



- Parses <u>SQL statements</u> for tables and columns which must be accessed
- Queries <u>statistics</u> from DB2 Catalog (populated by RUNSTATS utility)
- Determines least expensive <u>access path</u>
- Since it is a Cost-Based Optimizer it <u>chooses</u> the lease expensive access path



Optimizer Access Path Selection

- 1. Gets the current statistics from DB2 catalog for the columns and tables identified in the SQL statements. These statistics are populated by the Runstats utility.
- 2. Computes the estimated percentage of qualified rows for each predicate - which becomes the **filter factor** for the predicate.





Will a Scan or an Index Be Used?

A tablespace Scan sequentially reads <u>all</u> of the tablespace pages for the table being accessed.

Most of the time, the fastest way to access DB2 data is with an Index. For DB2 to consider using an index - the following criteria must be met:

- At least one of the predicates for the SQL statement must be indexable.
- One of the columns (in any indexable predicate) must exist as a column in an available index.

Will a Scan or an Index Be Used?

An index will **not** be used in these circumstances:

- When no indexes exist for the table and columns being accessed
- When the optimizer determines that the query can be executed more efficiently without using an index -
 - the table has a small number of rows or
 - using the existing indexes might require additional I/O - based on the cardinality of the index and the cluster ratio of the index.

Types of Indexed Access

Direct Index Lookup

values must be provided for each column in the index

Matching Index Scan (absolute positioning)

can be used if the high order column (first column) of an index key is provided

Nonmatching Index Scan (relative positioning)

- can be used if the first column of the index is not provided
- can be used for non-clustered indexes
- can be used to maintain data in a particular order to satisfy the ORDER BY or GROUP BY

Index Only Access

can be used with if a value is supplied for all index columns - avoids reading data pages completely

Sequential Prefetch

- A read-ahead mechanism invoked to prefill DB2's buffers so that data is already in memory before it is requested. Can be requested by DB2 under any of these circumstances:
 - A tablespace scan of more than one page
 - An index scan in which the data is clustered and DB2 determines that eight or more pages must be accessed.
 - An index-only scan in which DB2 estimates that eight or more leaf pages must be accessed.







Database Services Address Space

- When an SQL statement requesting a set of columns and rows is passed to the RDS, the RDS determines the best mechanism for satisfying the request. The RDS can parse an SQL statement and determine its needs.
- When the RDS receives an SQL statement, it performs these steps:
- 1. Checks authorization
- 2. Resolves data element names into internal identifiers
- 3. Checks the syntax of the SQL statement
- 4. Optimizes the SQL statement and generates an access path



Database Services Address Space

- The RDS then passes the optimized SQL statement to the Data Manager for further processing.
- The function of the DM is to lower the level of data that is being operated on. The DM analyzes the request for table rows or index rows of data and then calls the Buffer Manager to satisfy the request.
- The Buffer Manager accesses data for other DB2 components. It uses pools of memory set aside for the storage of frequently accessed data to create an efficient data access environment.

Database Services Address Space

- The BM determines if the data is in the bufferpool already. If so - the BM accesses the data and send it to the DM. If not - it calls the VSAM Media Manager to read and send back the data to the BM, so it can be sent to the DM.
- The DM receives the data and applies as many predicates as possible to reduce the answer set. Only Stage 1 predicates are applied in the DM.

Database Services Address Space

- Finally, the RDS receives the data from the DM. All Stage 2 predicates are applied, the necessary sorting is performed, and the results are returned to the requestor.
- Considering these steps, realize that Stage 1 predicates are more efficient because they are evaluated earlier in the process, by the DM instead of the RDS, and thereby reduce overhead during the processing steps.

SQL Coding Strategies and Guidelines

- Understand Stage 1 and Stage 2 Predicates
- Tune the queries that are executed more frequently first!
- It Depends!
- Know Your Data!
- Static vs. Dynamic SQL
- Batch vs. Interactive (CICS vs. web)



Rows Returned

- Minimize the number of rows searched and/or returned
- Code predicates to limit the result to only the rows needed
- Avoid generic queries that do not have a WHERE clause







Use For Fetch Only

- When a SELECT statement is used only for data retrieval - use FOR FETCH ONLY
- FOR READ ONLY clause provides the same function and is ODBC compliant
- Enables DB2 to use 'block fetch'
- Monitor the performance to decide which is best for each situation



<section-header><section-header><section-header><list-item><list-item><list-item><list-item><list-item><list-item><list-item>



Use Inline Views

- Inline views allow the FROM clause of a SELECT statement to contain another SELECT statement
- May enhance performance of the outer select by applying predicates in the inner select
- Useful when detail and aggregated data must be returned in a single query



Avoid Data Conversion When comparing column values to host variables - use the same Data Type Length When DB2 must convert data, available indexes are sometimes <u>not used</u>





Example: Outer Join With A Local Predicate

SELECT emp.empno, emp.lastname, dept.deptname FROM emp LEFT OUTER JOIN dept ON emp.workdept = dept.deptno WHERE emp.salary > 50000.00;

Works correctly but... the outer join is performed first, before any rows are filtered out.

Example: Outer Join Using An Inline View

SELECT emp.empno, emp.lastname, dept.deptname FROM (SELECT empno, lastname

FROM emp WHERE salary > 50000.00) as e LEFT OUTER JOIN dept ON emp.workdept = dept.deptno

Works better... <u>applies the inner join predicates first</u>, reducing number of rows to be joined



Use BETWEEN



Except when comparing a host variable to 2 columns

Stage 2 : WHERE

:hostvar BETWEEN col1 and col2

Stage 1: WHERE

Col1 <= :hostvar AND col2 >= :hostvar



Use LIKE With Care

- Avoid the % or the _ at the beginning because it prevents DB2 from using a matching index and may cause a scan
- Use the % or the _ at the end to encourage index usage



Use EXISTS

Use EXISTS to test for a condition and get a True or False returned by DB2 and not return any rows to the query:

SELECT col1 FROM table1 WHERE EXISTS (SELECT 1 FROM table2 WHERE table2.col2 = table1.col1)



- After the indexes, place the predicate that will eliminate the greatest number of rows first
- Know your data
 - Race, Gender, Type of Student, Year, Term





Other Cautions

- Predicates that contain concatenated columns are not indexable
- SELECT Count(*) can be expensive
- CASE Statement powerful but can be expensive

With OPTIMIZE for n ROWS

- For online applications, use 'With OPTIMIZE for n ROWS' to attempt to influence the access path DB2 chooses
- Without this clause, DB2 chooses the best access path for <u>batch processing</u>
- With this clause, DB2 optimizes for quicker response for online processing
- Try Optimize for 1, for 10, for 100

Review DB2 Optimizer

- DB2 is a Cost-based optimizer
- RUNSTATS populates the DB2 Catalog
- DB2 Catalog used to determine access path
- Create Indexes for columns you frequently select and sort

- Avoid Unnecessary Sorts in SQL
- Code the SQL predicates thoughtfully



Filter Factors for Predicates

- Filter factor is based on the <u>number of rows</u> that will be filtered out by the predicate
- A ratio that estimates I/O costs
- The lower the filter factor, the lower the cost, and in general, the more efficient the query

Review the handout as we discuss this topic

Filter Factors for DB2 Predicates Filter Factor Formulas - use FIRSTKEYCARDF column from the SYSINDEXES table of the Catalog If there are no statistics for the indexes, the default filter factors are used The lowest default filter factor is .01 : Column BETWEEN Value1 AND Value2 Column LIKE 'char%' Equality predicates have a default filter factor of .04 : Column = :hostvalue Column = :hostvalue ColumnA = ColumnB (of different tables)













Review Filter Factors for Predicates

- DB2 Catalog
- Filter Factors
- Column Matching
- Order of Predicate Evaluation

Runstats and Reorg

Runstats Utility

- updates the catalog tables with information about the tables in your system
- used by the Optimizer for determining the best access path for a SQL statement
- Reorg Utility
 - reorganizes the data in your tables
 - good to run the RUNSTATS after a table has been reorg'd
- Use Workbench to review the statistics in both Test and Production databases





DB2 Explain

- Required and reviewed by DBA when a DB2 program is moved to production
- Recognizes the <u>? Parameter Marker</u> assumes same data type and length as you will define in your program
- Know your data, your table design, your indexes to maximize performance







DB2 Explain Columns

- QUERY Number Identifies the SQL statement in the PLAN_TABLE (any number you assign - the example uses the numeric part of the userid)
- BLOCK query block within the query number, where 1 is the top level SELECT. Subselects, unions, materialized views, and nested table expressions will show multiple query blocks. Each QBLOCK has it's own access path.
- PLAN indicates the order in which the tables will be accessed



DB2 Explain Columns METHOD - shows which JOIN technique was used: O0- First table accessed, continuation of previous table accessed, or not used.

- 01- Nested Loop Join. For each row of the present composite table, matching rows of a new table are found and joined
- 02- Merge Scan Join. The present composite table and the new table are scanned in the order of the join columns, and matching rows are joined.
- 03- Sorts needed by ORDER BY, GROUP BY, SELECT DISTINCT, UNION, a quantified predicate, or an IN predicate. This step does not access a new table.
- 04- Hybrid Join. The current composite table is scanned in the order of the join-column rows of the new table. The new table accessed using list prefetch.
- **TNAME** name of the table whose access this row refers to. Either a table in the FROM clause, or a materialized VIEW name.
- **TABNO** the original position of the table name in the FROM clause









DB2 Explain Analysis

Guidelines:

- You want to avoid tablespace scans (TYPE = R) or at least be able to explain why. Tablespace scans are acceptable for small tables.
- Nested Loop Join is usually the most efficient join method.
- Index only access is desirable (but usually not possible)
- You should strive for Index access with the matching columns being the same as the number of columns in the index.

DB2 Explain Analysis

Try to answer the following questions:

- Is Access through an Index? (TYPE is I, I1, N or MX)
- Is Access through More than one Index (TYPE is M, MX, MI or MU)
- How many columns of the index are used in matching (TYPE is I, I1, N or MX and MC contains number of matching columns)
- Is the query satisfied using only the index? (IO = Y)





DB2 Explain Example - 5 tables

SELECT C.COURSE_NUMBER, C.COURSE_IND, C.YEAR, C.TERM, C.SECTION_NUMBER, C.SUMMER_SESSION_IND, C.FACULTY_ID, E.COURSE_DEPT_NUMBER,
D.LAST_NAME AS FACULTY_LAST_NAME,
D.FIRST_NAME AS FACULTY_FIRST_NAME,
D.MIDDLE_NAME AS FACULTY_MID_NAME,
A.STUDENT_ID, A.HOURS,
B.LAST_NAME AS STUDENT_LAST_NAME,
B.FIRST_NAME AS STUDENT_FIRST_NAME,
B.MIDDLE_NAME AS STUDENT_MID_NAME,
B.CURR_CLASS, B.CURR_DIV, B.CURR_MAJOR, B.RACE, B.GENDER



DB2 Explain Example - 5 tables FROM FSDBA.COURSE MASTER AS E, FSDBA.CURRENT COURSES AS C, FSDBA.TEACHER MASTER AS D, FSDBA.STUDENT COURSE AS A, FSDBA.STUDENT MASTER AS B WHERE C.COURSE NUMBER = E.COURSE NUMBER AND C.COURSE IND = E.COURSE IND ANDC.FACULTY ID = D.FACULTY ID AND C.YEAR = A.YEAR AND C.TERM = A.TERMAND C.COURSE NUMBER = A.COURSE NUMBER AND C.COURSE IND = A.COURSE IND ANDC.SECTION NUMBER = A.SECTION NUMBER AND A.STUDENT ID = B.STUDENT ID AND 76





DB2 Explain Example - 5 tables

ORDER,BY C.COURSE_NUMBER, C.COURSE_IND, C.SECTION_NUMBER, STUDENT_LAST_NAME, STUDENT_FIRST_NAME, STUDENT_MID_NAME FOR FETCH ONLY OPTIMIZE FOR 15 ROWS;

DB2 Explain Example - 5 tables

+	+		+	+++	+			
QUERY	BLOCK	PLAN	METH	TNAME	TABNO	TYPE	MC	ANAME
+	+-	+		++++++	+			
00587	01	01	00	CURRENT_COURSES	02	R	00	
00587	01	02	04	COURSE_MASTER	01	Ι	02	IXCRM01
00587	01	03	04	STUDENT_COURSE	04	Ι	03	IXSTC02
00587	01	04	01	TEACHER_MASTER	03	Ι	01	IXTCM01
00587	01	05	01	STUDENT_MASTER	05	Ι	01	IXSTM01
00587	01	06	03	-	00		00	



SET CURRENT SQLID='FSUDBA'; EXPLAIN PLAN SET QUERYNO=587 FOR SELECT MON, SUM(AMOUNT) FROM (SELECT) MACH_DATE, MONTH(MACH_DATE) AS MON, SUM (AMOUNT) AS AMOUNT FROM FSUDWH.SAMAS_TRANSACTIONS SAM, FSUDWH.FUND_CODES FND, FSUDWH.OBJECT_CODES OBJ, FSUDWH.APPRO_CATEGORY_CDS CAT

Example: SAMAS Query1

WHERE (CAT.APPRO_CATEGORY= SAM.APPRO_CATEGORY) AND (OBJ.OBJECT_CODE= SAM.CHARGE_OBJECT) AND (SAM.STATE_FUND=FND.STATE_FUND AND SAM.FUND_ID=FND.FUND_CODE) AND ((SAM.RECORD_TYPE = 'I') AND SAM.CHARGE_ORG LIKE '021000000' AND SAM.MACH_DATE BETWEEN '2000-07-01' AND '2001-06-30' AND (SAM.B_D_E_R = 'D' AND SAM.TRANS_TYPE <> '80' AND SAM.RECORD_TYPE = 'I')) GROUP BY SAM.MACH_DATE) AS QRY1 GROUP BY MON ; 83



Example: SAMAS Query1

QUERY	BLOCK	PLAN	METH	TNAME	TABNO	TYPE	MC	ANAME
00587	01	01	00	QRY1	01	R	00	
00587	01	02	03		00		00	
00587	02	01	00	SAMAS_TRANSACTIONS	02	I	01	IXSTR08
00587	02	02	01	FUND_CODES	03	I	02	IXFUN01
00587	02	03	01	OBJECT_CODES	04	I	01	IXOBJ01
00587	02	04	01	APPRO_CATEGORY_CDS	05	I	01	IXACC01
00587	02	05	03		00		00	
_	_	_	_		_	_	-	_
								85

Example: SAMAS Query1

IO	SNU	SNJ	SNO	SNG	SCU	SCJ	SCO	SCG	ΡĒ
Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	S
Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Y	
Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	S
Y	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	
Y	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	
Y	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	
Ν	Ν	Ν	Ν	Ν	Ν	Ν	Ν	Y	

Example: SAMAS Query2 SET CURRENT SQLID = 'FSUDBA'; EXPLAIN PLAN SET QUERYNO = 1 FOR SELECT MACH_DATE, FISCAL_YEAR, CHARGE_ORG, PRIMARY_DOC_NUM, AMOUNT FROM FSDBA.SAMAS_TRANSACTIONS WHERE MACH_DATE >= '2001-07-01'

AND MACH_DATE <= '2002-06-30' AND FISCAL_YEAR IN ('20012002' ,'20012002') AND BUDGET_ENTITY = '48900100'





			Ех	car	nŗ	ole	:: {	SA	M	AS	Qı	ıer	y2	2	
	QUERY BLOCK PLAN METH TNAME									TABNO	TYPE	MC	ANAME		
	00001 01 01 00001 01 02			01 02	0(0 SAMAS_TRANSACTIONS					01 00	R	00 00		
1	I0 	SNU	SNJ	SNO	SNG	SCU	SCJ	SCO	SCG	PF 					
н	N N	N N	N N	N N	N N	N N	N N	N Y	N N	S					
н															
н															



