Performance Hero
Success Guide

Brother-Panther™ for DB2 LUW

Scott Hayes
President & CEO, DBI
IBM DB2 GOLD Consultant
Scott_Hayes@Database-Brothers.com

www.database-brothers.com
Phone: 512-249-2324

This page intentionally left blank.

www.database-brothers.com
Phone: 512-249-2324

PATIENCE IS A VIRTUE

Fully Optimized Performance Tuning requires *Patience* and *Discipline*

This paper describes the discipline. We cannot teach you patience, but we can hope you like to take coffee and lunch breaks.

www.database-brothers.com
Phone: 512-249-2324

Introduction

This paper is not about the intimate details of product installation and configuration. It is assumed you have already read the product installation and configuration guide, successfully installed Brother-Panther's components, and possibly started collection for one or more databases.

The purpose of this paper is to help you become a Performance Hero in your organization relatively quickly. If you follow the steps and advice herein, it is highly likely that you will be very successful in achieving remarkable, measurable performance improvements in a short amount of time.

That being said, *patience is required*. In order to be extremely successful in database tuning, you must understand the database Statement workload and its aggregate costs. Workloads do not run instantly and they cannot be determined from snapshots or small windows of time. Workloads run over an extended period of time. The better you know the database Statement workload over an extended period of time, the more accurately you will be able to understand the workload's costs, and the more successful you will be at reducing costs of Statement execution. When you successfully reduce costs, database performance will improve resulting in faster query times, reduced CPU utilization, the ability to accommodate more users on the same hardware, or improved query throughput.

Brother-Panther versus Brother-Eagle™

Brother-Panther is designed to analyze database workloads over time. If you want or need to know "What's happening right now?!?!?" then Brother-Eagle is the tool you should work with, for Brother-Eagle provides near real-time analysis of database performance metrics and drill downs to current activity and locks.

Brother-Panther's Collection Discipline

Brother-Panther works on the concept of collection windows or "intervals". This is done for a few reasons:

- Minimize the overhead of monitoring
- Manage the size of the performance repository database
- Enforce discipline in the methods, management, and analysis of activity and history

Collection Intervals occur by default on 15 minute time boundaries beginning at the top of the hour, quarter past the hour, half past the hour, and three quarters past the hour.

Performance information is saved to the repository database approximately at the end of each 15 minute interval. Collection does not begin until the start time of a collection interval.

By way of example, then, let us assume that you use the DBI Admin Console to start collection for a database at 8:50am. Actual collection will begin at the top of the hour, or 9:00am. The first performance information will be loaded into the repository database around 9:15am. Therefore, the Brother-Panther Console will have its first information to present to you shortly after the first full collection interval, or around 9:16am or shortly thereafter.

If you selected Autonomic monitoring of Statement activity when you started database collection with the DBI Admin Console, the autonomic sensors begin checking key metrics during the first and subsequent collection intervals. It is possible, then, that Autonomic Statement collection may not begin until 9:15am, and the first Statement Performance would be available for viewing after 9:30am.

OLTP Database Collection Recommendation

For best results, the incredibly successful DBA should ideally allow collection to occur for a given OLTP database for at least an hour before diving in and attempting analysis and changes. This is not to say that an hour is required, but it is recommended.

Data Warehouse Database Collection Recommendation

For Data Warehouse databases, the best tuning results will be obtained after collecting performance data for several hours or a few days. Remarkably, even in seemingly Adhoc database environments, patterns of repetitious Statements often emerge as end users execute the same queries daily or weekly with same or different values in search arguments. It is for this reason that Brother-Panther provides the capability to analyze and aggregate Statement activity and costs over extended periods of time, such as an entire month, so that you can determine repetitious patterns of costly activity.

With this introduction behind us, and the mutual understanding of the importance and value of patience, we'll next take a look at the disciplined steps to becoming a Performance Hero in your organization.

Step 1: Start Collection

Use the DBI Admin Console to Start Collection on databases that are important to your organization. You can add several different databases hosted on the same or different database servers if you like. All databases that you monitor will appear in Brother-Panther's Console "Database Score" grid, which will allow you to see, at a glance, which databases have the most "opportunity for improvement" (that is, the lowest score, or the most problems).

DBI understands that some DBAs may not be interested in the performance of all databases, so each Brother-Panther user can set filter criteria to omit databases from the display which are not of interest.

After you have started collection, the Databases view of your DBI Admin Console might look similar to this:

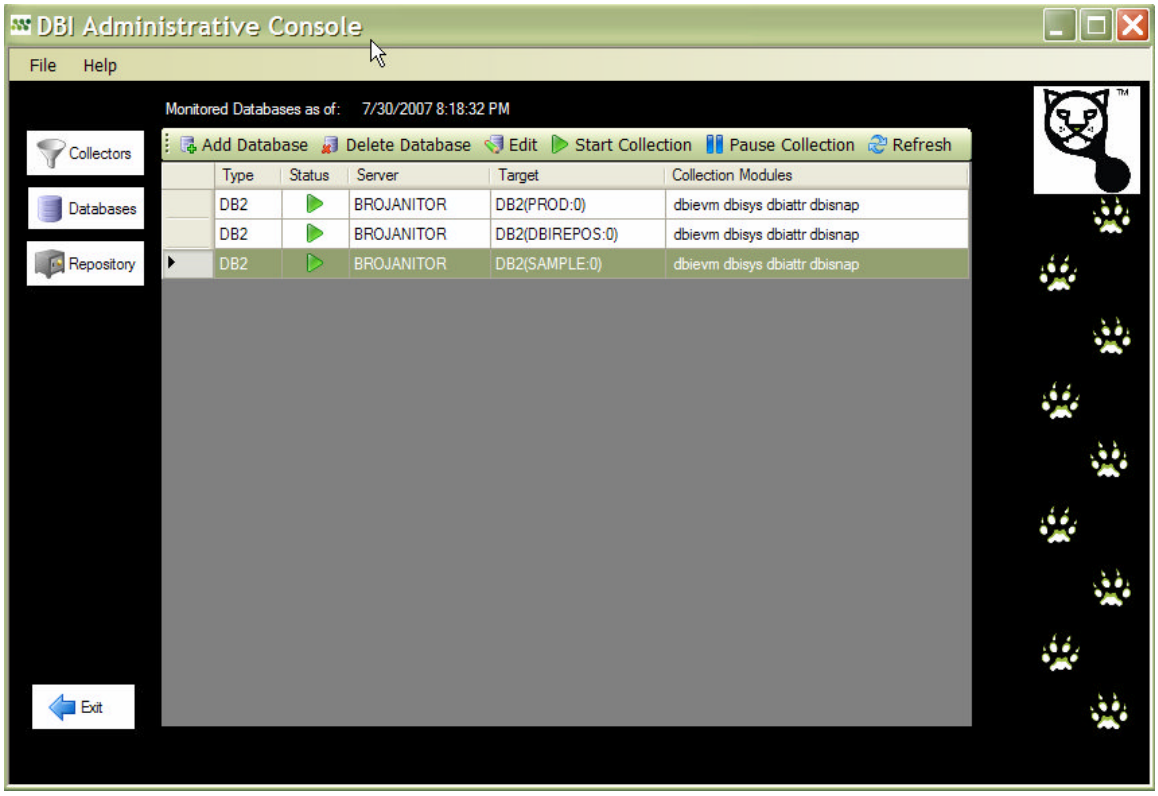

You will note that 3 databases are being monitored on server BROJANITOR, instance DB2. One of the 3 databases is the DBI Repository. As the TV commercials say, "Do Not Attempt this at Home"

Step 2: Take a Break

Get coffee.
Go to lunch.
Remember, patience is a virtue. The longer you collect performance data, the greater your chances for remarkable success.

Step 3: Start the Brother-Panther Console

Click the Windows Start button, choose All Programs, choose Database-Brothers, and then choose “Brother-Panther DB2”. The Database Score grid should appear as shown below:

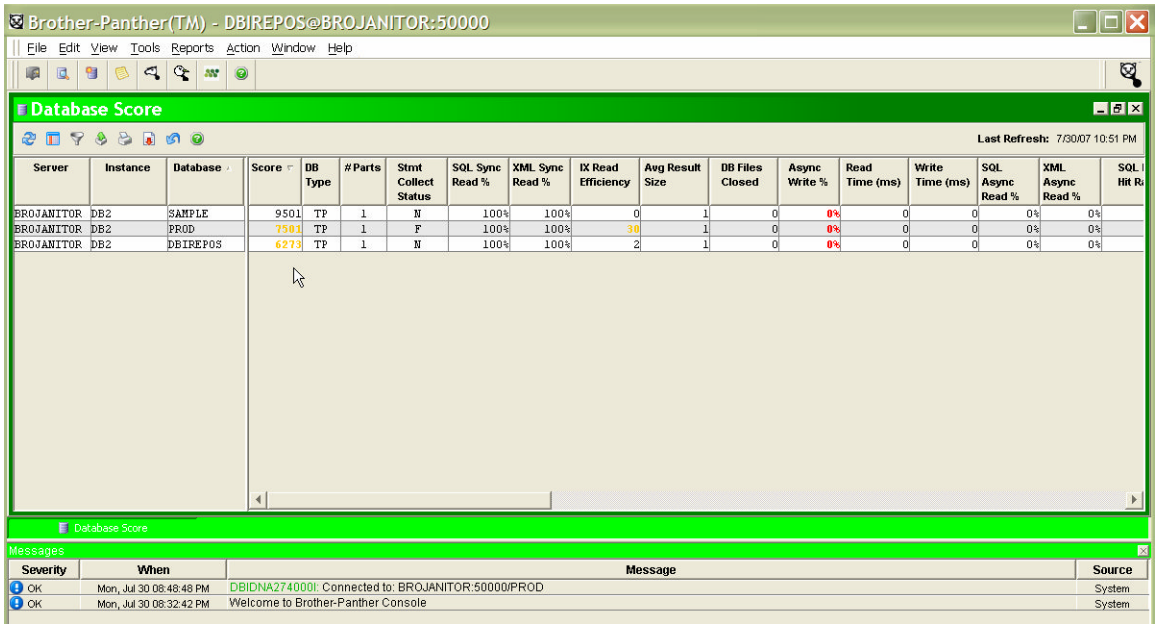

www.database-brothers.com
Phone: 512-249-2324

Step 4: Check your Database Scores

Click on the “Score” column heading until the database performance grid sorts the database scores from lowest to highest.

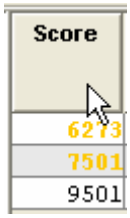

Databases with the lowest scores have the most detected ‘opportunities for improvement’. 10,000 is the maximum score.

Step 5: Check the Database Score Reports

Left click on the database name of a database having a score less than 10,000.

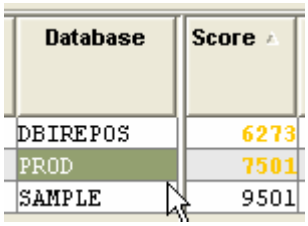

Next, right click on that database name and select the option to Analyze DB Score:

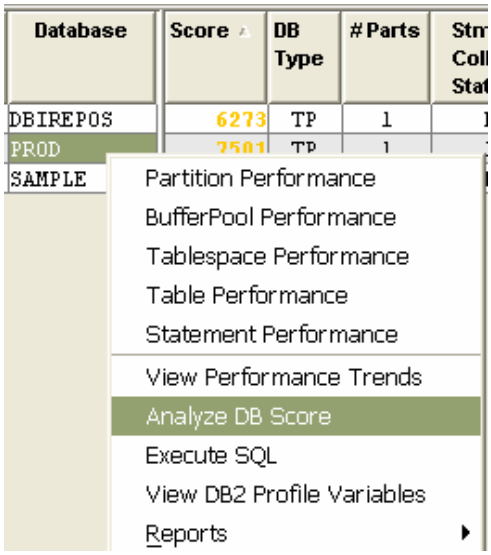

www.database-brothers.com
Phone: 512-249-2324

The Database Score report will describe the analysis and factors that contributed to the Database Score. Some, but not all, Analysis messages are accompanied by Recommendations. Carefully review the Analysis and any Recommendations. The Database Score Report may provide you with advice or tips that will point you towards changes that will produce immediate, substantial, and measurable value.

For this particular database, the ratio of Rows Read to Rows Fetched (Index Read Efficiency, or IREF for short) is unreasonably high compared to industry expectations and best practices. The Recommendation suggests that indexes may be missing or sub-optimally defined.

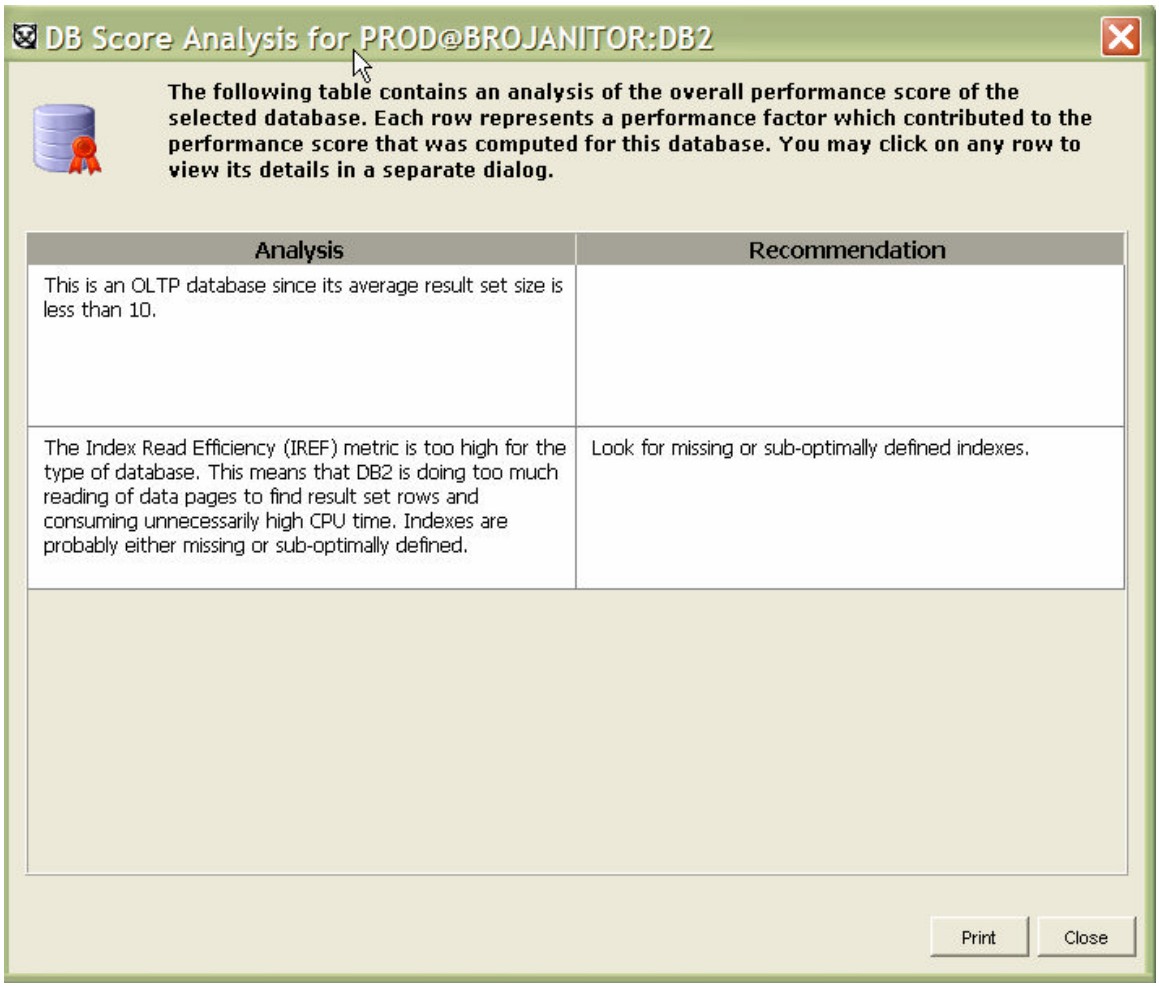

To learn more about IREF, consult the Brother-Panther documentation or Brother-Eagle's Expert Advice at http://www.database-brothers.com/brother-eagle/advice/db2dbiref.php.

www.database-brothers.com
Phone: 512-249-2324

Here is another sample Database Score Report:

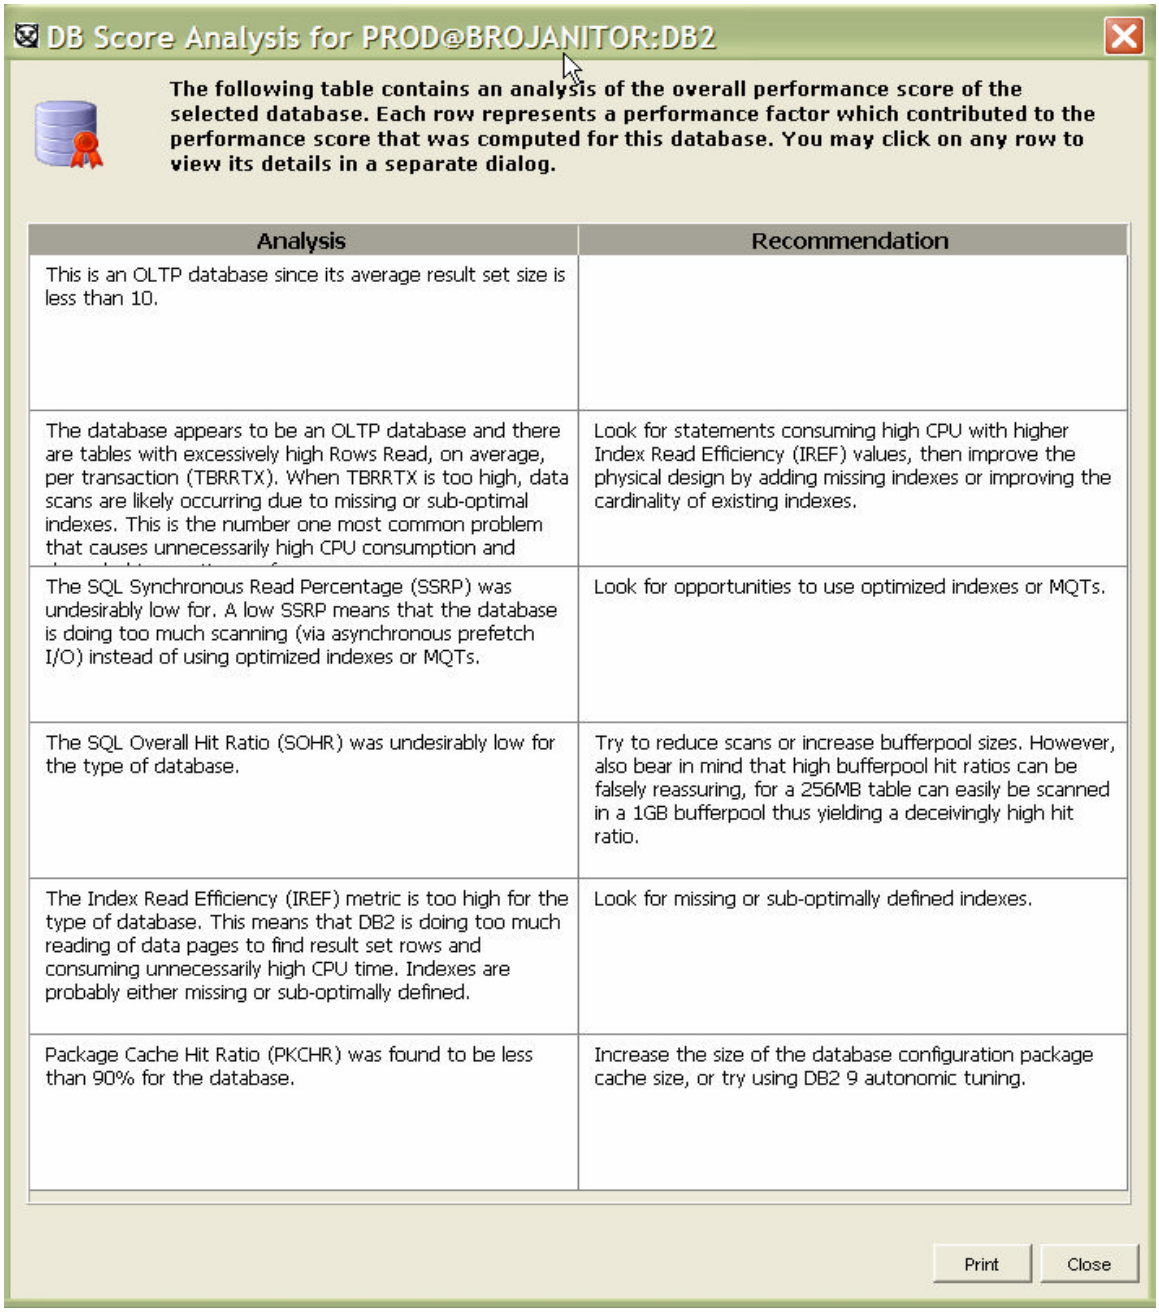

www.database-brothers.com
Phone: 512-249-2324

Step 5: Check Table Performance

The most typical and fastest route to value is an examination of Table Performance. Table Performance will show us tables with the highest Rows Read or Rows Written, along with other potential problems such as unusually high Rows Read per Transaction or Overflows.

Right click on the database name again (PROD), and then select the option Table Performance:

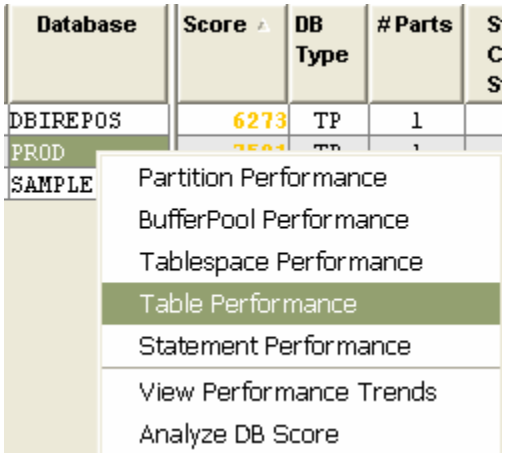

You will see a Table Performance grid like this example:

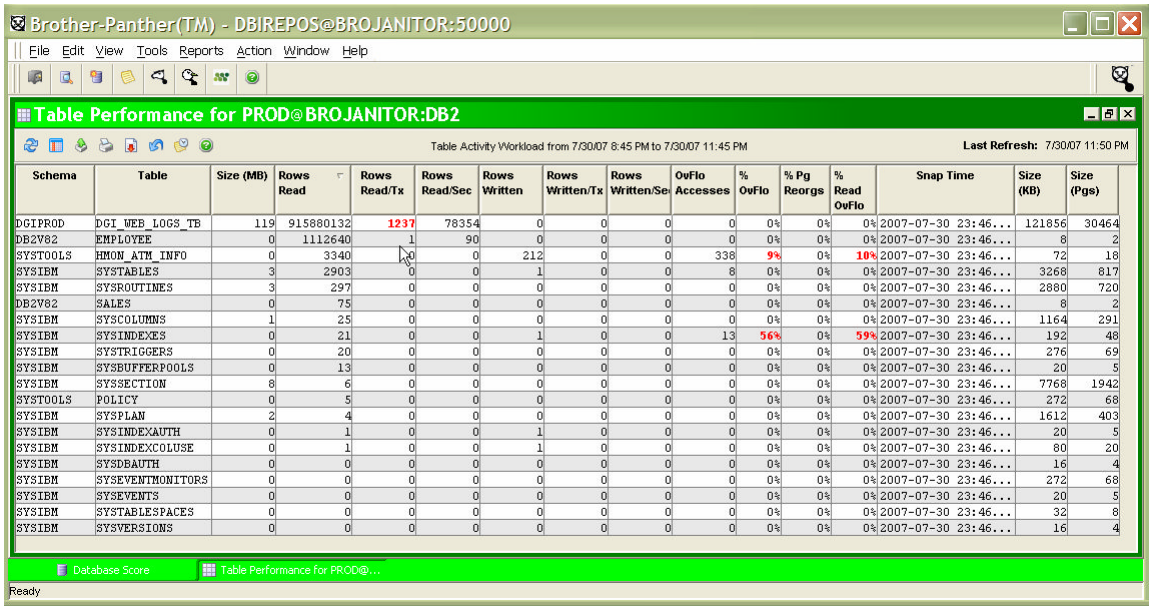

Schema	Table	Size (MB)	Rows Read	Rows Read/Tx	Rows Read/Sec	Rows Written	Rows Written/Tx	Rows Written/Se	OvFlo Accesses	% OvFlo	% Pg Reorgs	% Read OvFlo	Snap Time	Size (KB)	Size (Pgs)
DGIPROD	DGI_WEB_LOGS_TB	119	915880132	1237	78354	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	121856	30464
DB2V82	EMPLOYEE	0	1112640	1	90	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	8	2
SYSTOOLS	HMON_ATM_INFO	0	3340	0	0	212	0	0	338	9%	0%	10%	2007-07-30 23:46...	72	18
SYSIBM	SYSTABLES	3	2903	0	0	1	0	0	8	0%	0%	0%	2007-07-30 23:46...	3268	817
SYSIBM	SYSROUTINES	3	297	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	2880	720
DB2V82	SALES	0	75	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	8	2
SYSIBM	SYSCOLUMNS	1	25	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	1164	291
SYSIBM	SYSINDEXES	0	21	0	0	1	0	0	13	56%	0%	59%	2007-07-30 23:46...	192	48
SYSIBM	SYSTRIGGERS	0	20	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	276	69
SYSIBM	SYSBUFFERPOOLS	0	13	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	20	5
SYSIBM	SYSSECTION	8	6	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	7768	1942
SYSTOOLS	POLICY	0	5	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	272	68
SYSIBM	SYSPLAN	2	4	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	1612	403
SYSIBM	SYSINDEXAUTH	0	1	0	0	1	0	0	0	0%	0%	0%	2007-07-30 23:46...	20	5
SYSIBM	SYSINDEXCOLUSE	0	1	0	0	1	0	0	0	0%	0%	0%	2007-07-30 23:46...	80	20
SYSIBM	SYSDBAUTH	0	0	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	16	4
SYSIBM	SYSEVENTMONITORS	0	0	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	272	68
SYSIBM	SYSEVENTS	0	0	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	20	5
SYSIBM	SYSTABLESPACES	0	0	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	32	8
SYSIBM	SYSVERSIONS	0	0	0	0	0	0	0	0	0%	0%	0%	2007-07-30 23:46...	16	4

www.database-brothers.com
Phone: 512-249-2324

Click on the “Rows Read” column until the grid is sorted in descending sequence; the tables with the highest Rows Read will appear at the top.

Generally, Rows Read per Transaction shouldn’t exceed the value of 10. Here, the table DGIPROD.DGI_WEB_LOGS_TB has had over 915,000,000 Rows Read with Rows Read per Transaction of 1,217. When Rows Read per Transaction exceeds 10, it usually means that scans are occurring against the table data - which means that indexes are missing or sub-optimally defined.

Step 6: Find the Statement Activity causing High Rows Read

Right click on the table name of a table having the highest Rows Read or Rows Read/TX counts, and choose the option “Statement Performance”:

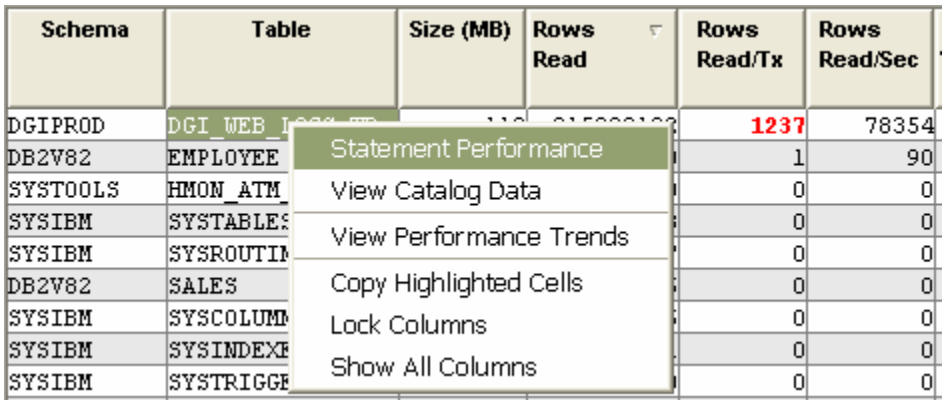

A Statement Performance grid will be displayed which shows the Statement activity that is driving I/O to the selected table. From the Table Performance grid, it is important to understand that Statement Performance shows all Statements driving I/O to the table, but other Statements are filtered out. This analysis is INVALUABLE to physical design and the determination of indexes, clustering indexes, MDC tables, and MQTs.

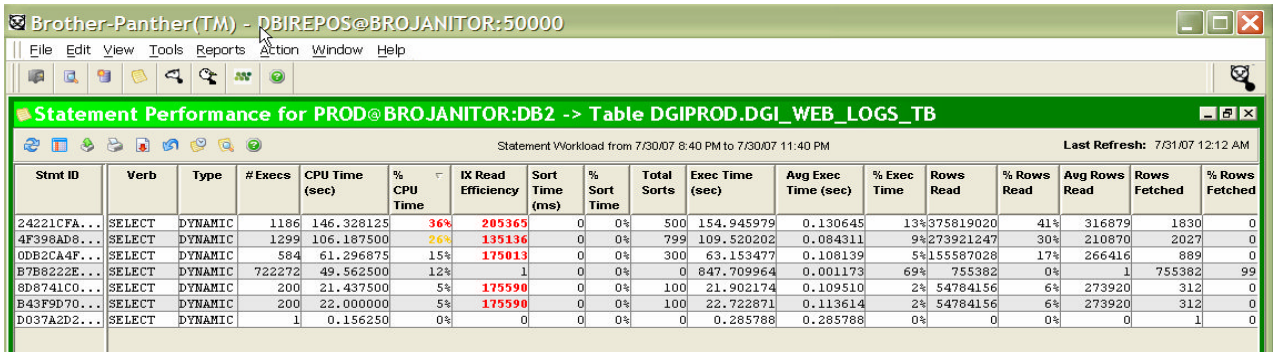

This Statement Performance grid shows some Statements executing that have horrific performance attributes. Index Read Efficiency (IREF) values are extremely high with the worst Statement having an IREF > 205,000! IREF > 10 is usually a problem, and IREF > 100 is ordinarily a performance crisis in the making.

Before moving on to the next step, please note that you can click on any column heading to sort the data by that performance metric. In this way, you can find the Statements having the highest aggregate CPU cost or Sort cost, the Statements having the highest average elapsed times, the highest/worst IREF, the highest Rows Read or Rows Written or percentages thereof, or the Statements that are retrieving (Rows Fetched) the highest quantity of Rows. And, remember, these are aggregated across matching Statement text patterns irrespective of any literal data values as described by US Patent #6,772,411.

Step 7: Reduce the Execution Cost of Expensive Statements

If a well stated problem is a half solved problem, at the top of your Statement Performance grid you should now have a problematic, costly Statement. To improve performance of the database, you must work to reduce the execution costs of the most costly Statements.

Step 7a: View the Statement text

Left click on the Verb column, or any column, of the grid Row that contains a costly Statement (often the top row), then Right click that column and select the option to "View Statement":

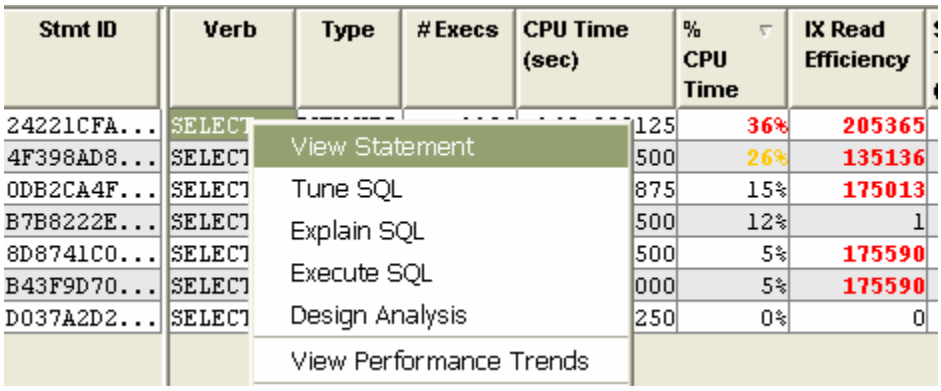

www.database-brothers.com
Phone: 512-249-2324

The Statement text will be shown in a pop-up window:

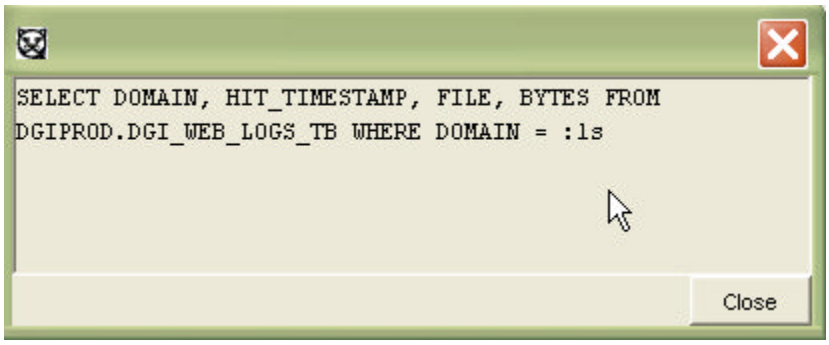

Statement text in this window can be copied and pasted into other windows such as Email or Word or Brother-Panther's Execute SQL feature.

Step 7b: Try the "Tune SQL" option

The Tune SQL option for a Statement will provide the Statement text, the properties of the Compilation environment, describe high cost steps of the Statement's internal execution, and show alternatives.

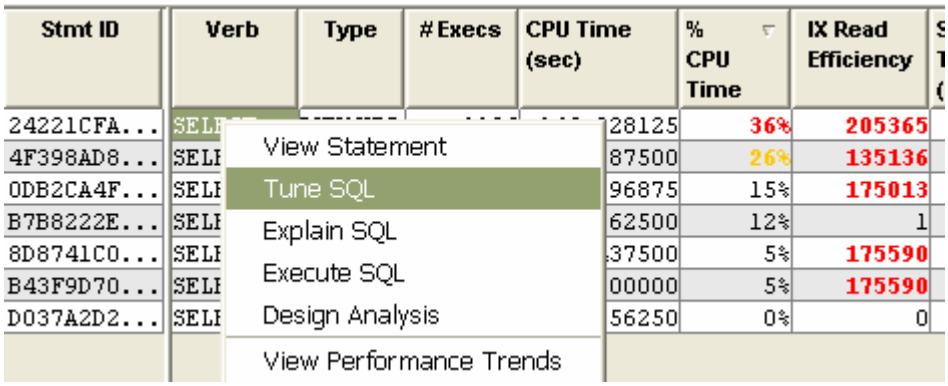

www.database-brothers.com
Phone: 512-249-2324

Sample Tune SQL display:

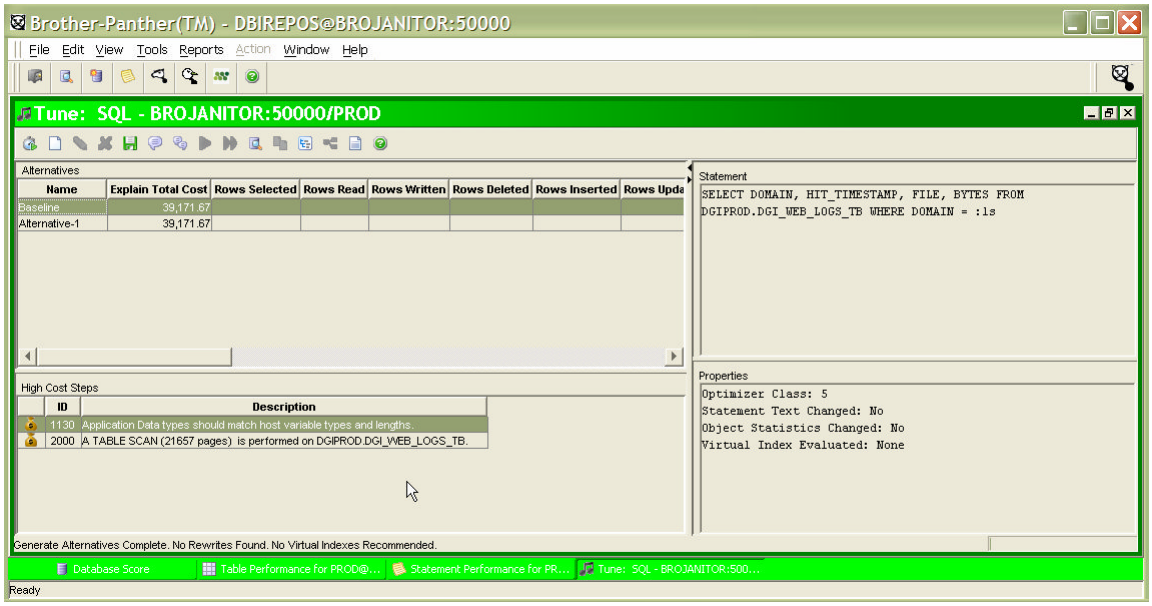

From this window, you can easily complete the next step which is to examine an Explain Plan View, or “Explain” the Statement. Click the Explain Plan View icon:

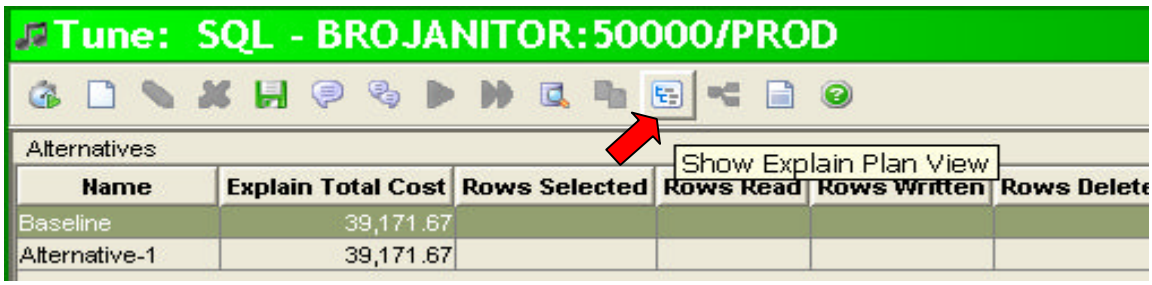

Also from this Tune SQL window you can easily obtain physical design analysis from the IBM Design Advisor:

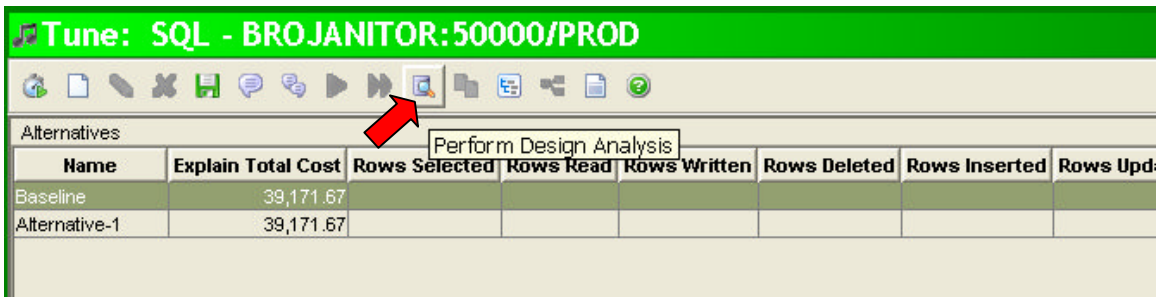

www.database-brothers.com
Phone: 512-249-2324

Step 7c: Review Explain

Again, from the Statement Performance grid, Right Click on a row having a disproportionately high execution cost, and select the option “Explain SQL” (You could, of course, have clicked the Explain icon in the previous step):

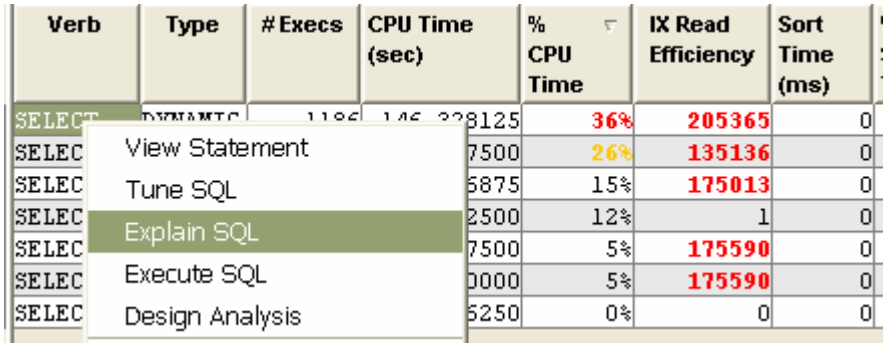

You will receive an Explain report formatted similar to this:

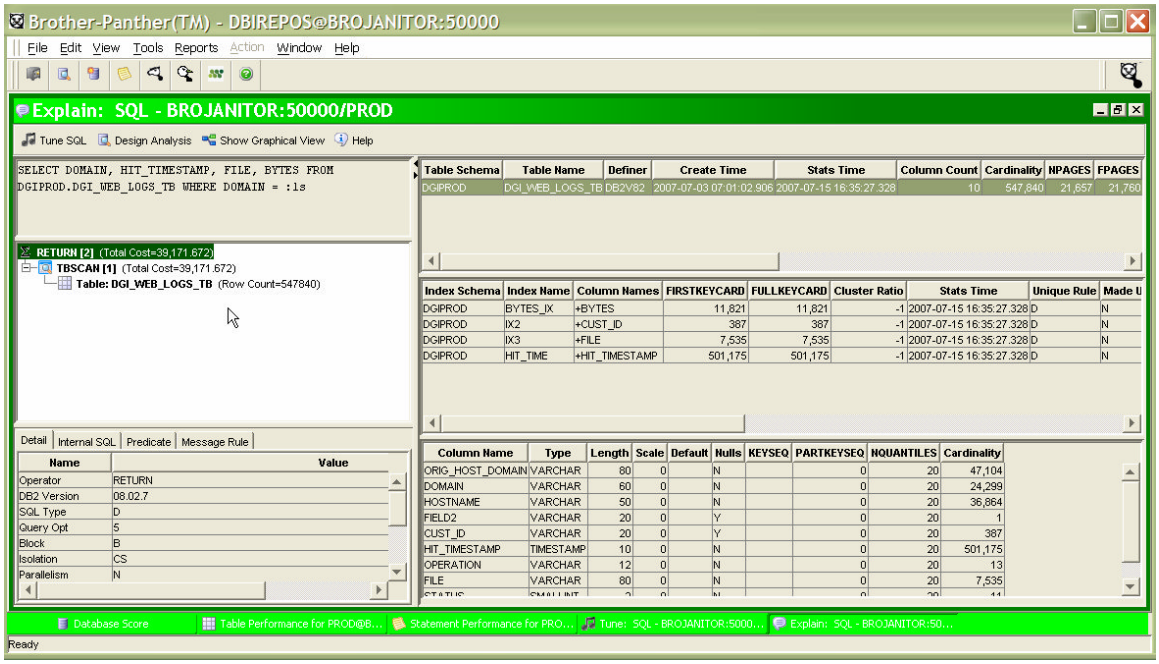

In one convenient window, you will find the Statement text, a graphical representation of the Explain plan, a summary of the compilation environment, statistics about the Table(s) involved, definitions and statistics for indexes on the table(s), and a description of the table definition, its columns, and their statistics. Talk about one-stop shopping. This Explain Analysis brings it all together.

www.database-brothers.com
Phone: 512-249-2324

Step 7d: Design Analysis

Again, from the Statement Performance grid, Right Click on a row having a disproportionately high execution cost, and select the option “Design Analysis” (You could, of course, have clicked the Design Analysis icon in step 7b):

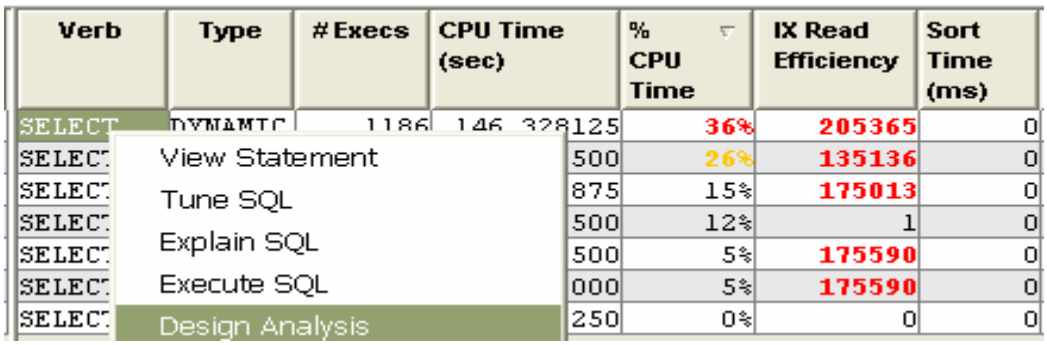

A Design Analysis option window will appear similar to this one:

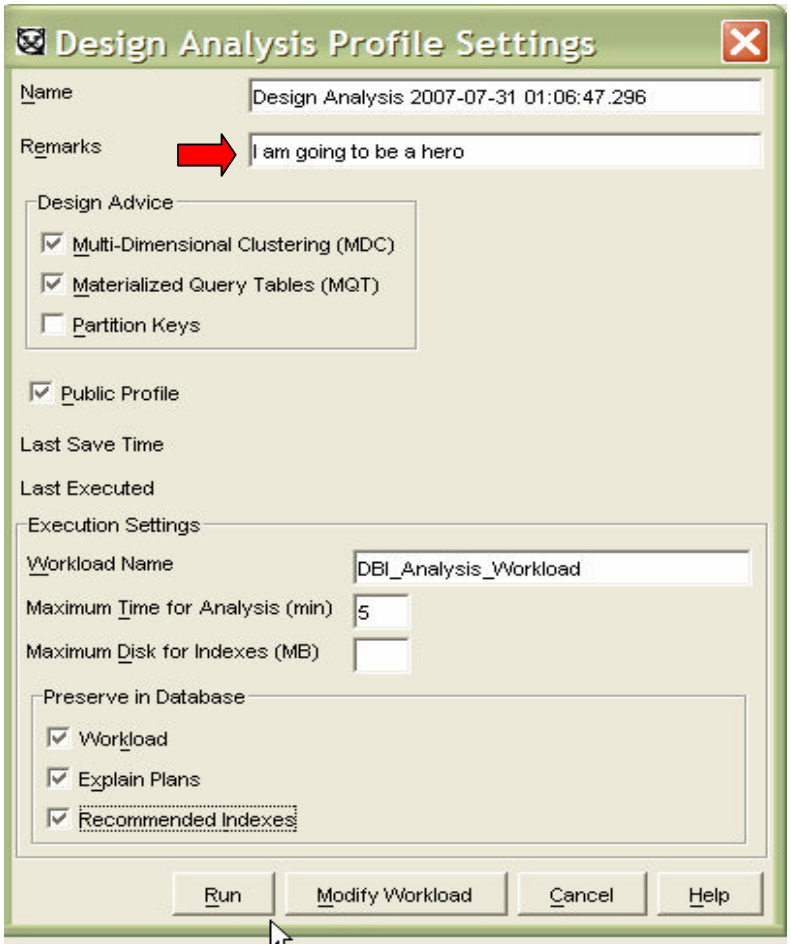

www.database-brothers.com
Phone: 512-249-2324

1. Add Remarks for documentation: “I am going to be a hero”
2. Check the MDC and MQT Design Advice boxes if you are seeking Design Advice for more than one statement (these boxes really shouldn’t be checked, but this is an example). Index advice will always be provided by default. You might select Partition Keys advice if your database has multiple partitions.
3. Give the Workload a Name: “DBI_Analysis_Workload”
4. Optionally provide a time limit for the analysis
5. Optionally provide a maximum amount of disk storage for proposed indexes (MB)
6. Check the Preserve in Database options as you see fit
7. **Click Run** to get Design Advice for the Statement selected earlier

A window similar to the one below will appear:

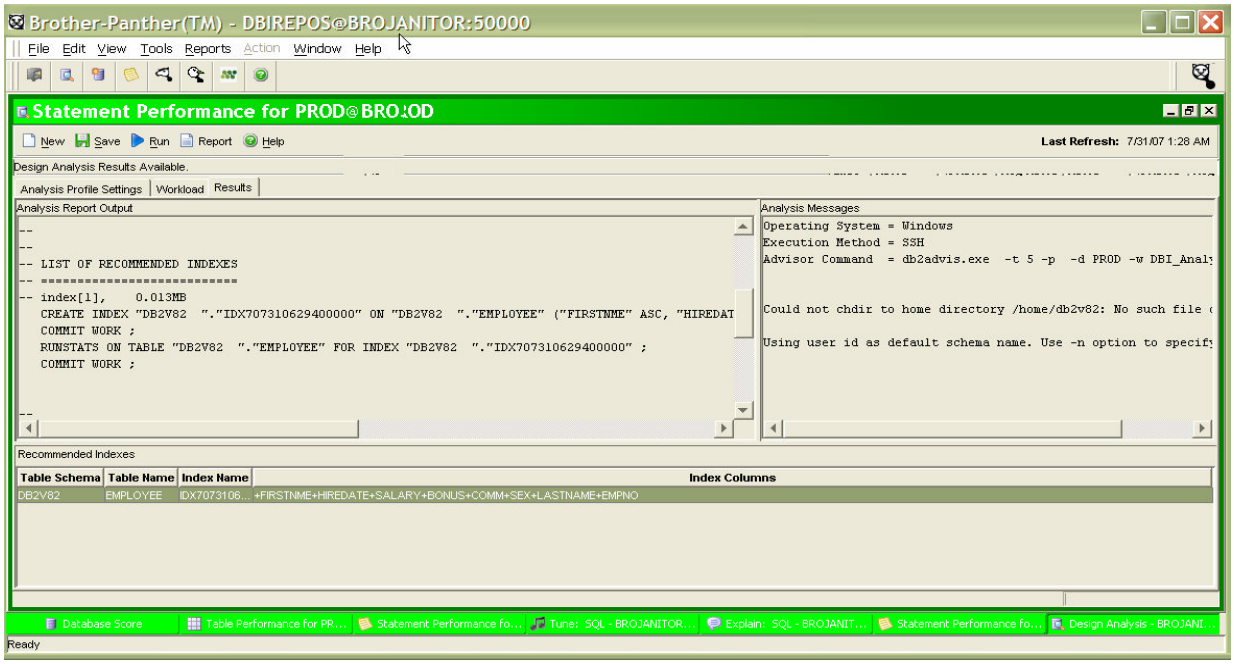

The top left pane contains the output from the IBM Design Advisor. The bottom pane contains a summary of proposed indexes.

The astute reader may notice that I have sought design advice for a different statement which accesses a different table for purposes of illustrating the procedure and Brother-Panther’s capabilities. You may also note that the IBM Design Advisor highly favors composite (multi-column) indexes so that Index Only Access can be achieved. The user is encouraged to exercise reasonable prudence when implementing Design Advisor recommendations - in other words, compare the WHERE predicates of the SQL to the columns in the proposed index; it may be appropriate to create a new index using only the first one or two columns of the Design Advisor’s recommendation.

Any text from the top left pane can be copied and pasted into Brother-Panther's "Execute SQL" function as an alternative method of creating new indexes or implementing other physical design changes. This alternate method would allow you to give any new indexes descriptive names of your choosing, rather than accepting the IBM generated gibberish names.

Step 7e: Implement Physical Design Improvement

From the Design Analysis window, you can easily create any recommended new index by Right Clicking on the index name and selecting the "Create" option. This will load and nicely format the CREATE INDEX statement into Brother-Panther's "Execute SQL" function. Here you can edit the CREATE INDEX command prior to running it. Be sure to:

1. Give the new Index a name that is descriptive and which conforms to your shop standards
2. Remove any extra or unnecessary columns from the index definition if the Design Advisor has become over zealous in its aggressive pursuit of Index Only access.

After your edits, if any, are complete, simply click the green "play" (right arrow) button at the top left to run the CREATE INDEX command (optionally followed by RUNSTATS)

.

Step 8: Take a Break

You deserve it.

Get some coffee or a brand name cola.

Take a peer to lunch and brag about how easy it was for you to find performance problems in the database, and fix them with just a couple mouse clicks.

Remember, patience is a virtue.

Step 9: Review Performance Trends

Return to the Database Score grid, Right Click on a Database name, and select the option "View Performance Trends":

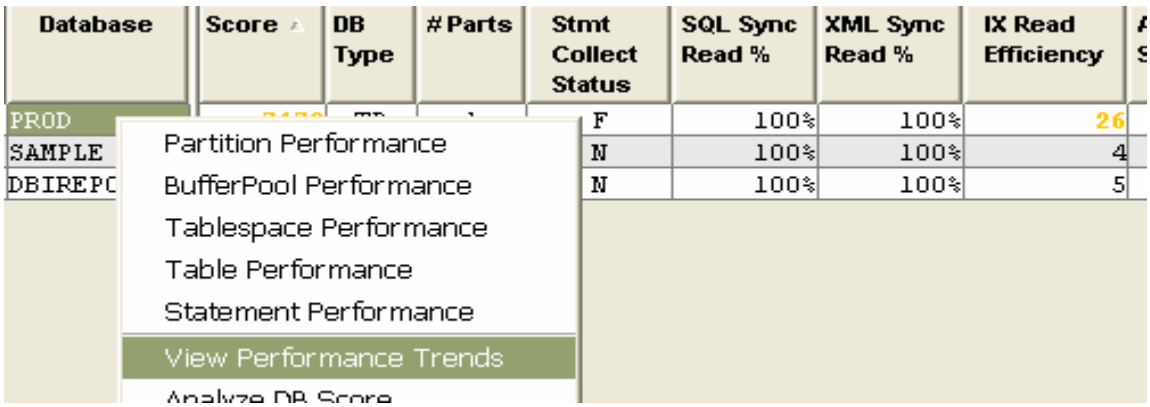

Another window will appear that provides you with several charting options. You can choose from bar or line charts, the time scale of the charts (hours, days, or months), the metrics to graph, and a time range for the graph:

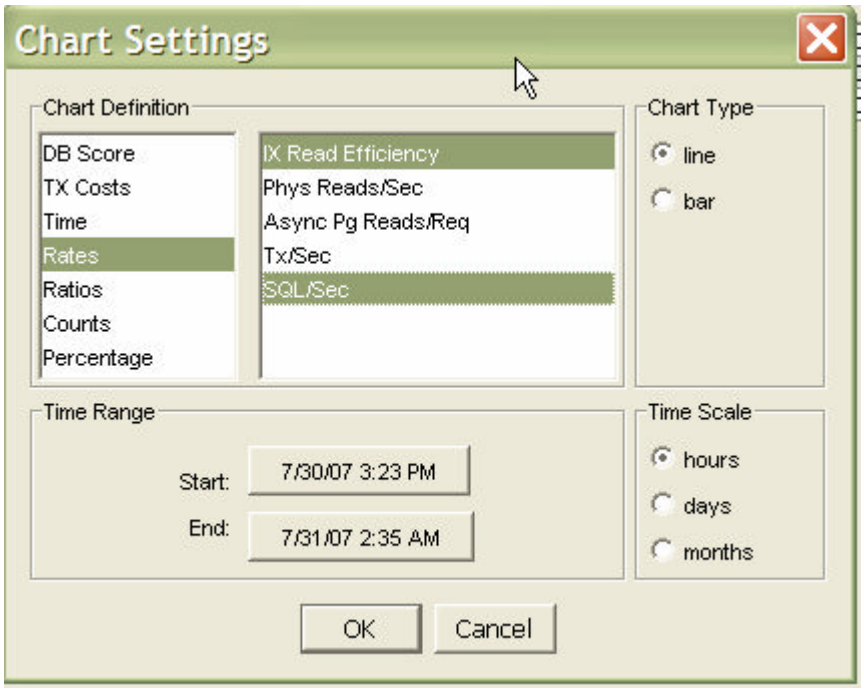

Click the OK button after you have made your trend graph option selections, and a graph window will appear:

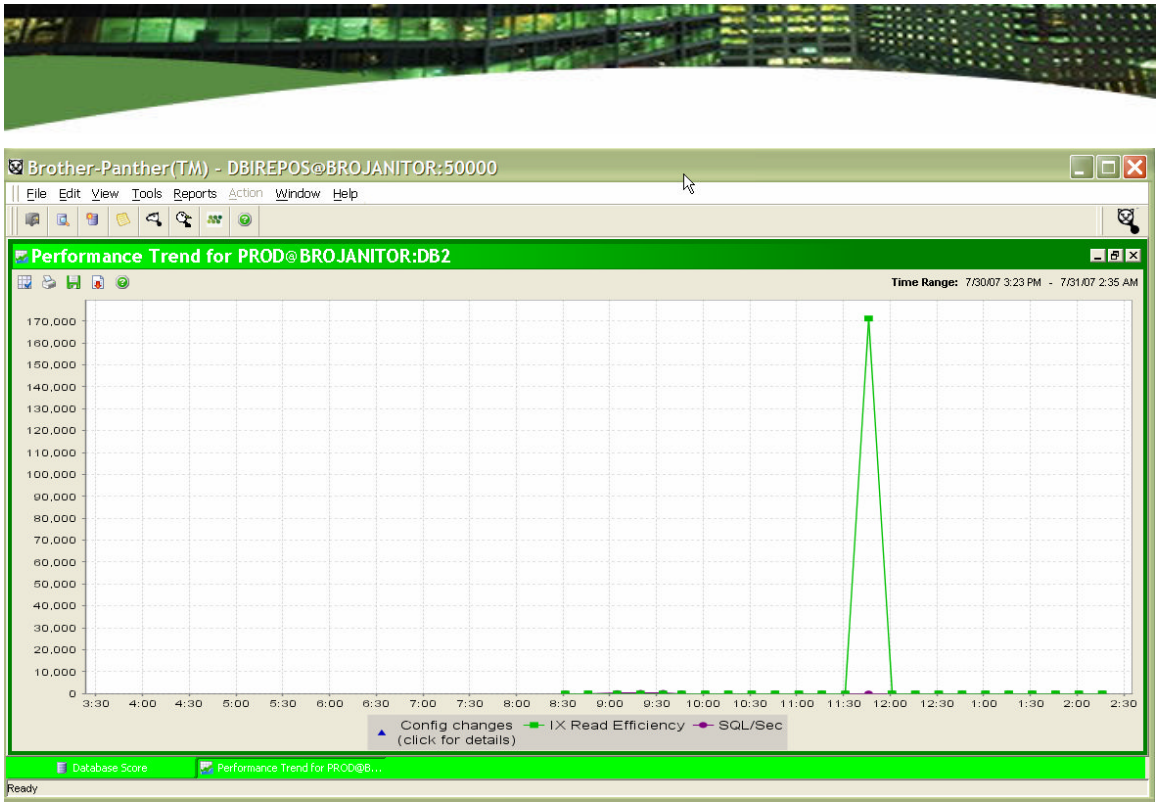

You will note that some metrics graph better together than others depending on the scale of values.

To verify the success of your tuning changes from Step 7, or any other configuration or physical design changes for that matter, some of the best metrics to graph include:

Category	Metric	Desired Movement
DB Score	Score	Higher is Better
TX Costs	Sorts/TX	Lower is Better
TX Costs	Hash Joins/TX	Lower is Better
TX Costs	Logical IX Reads/TX	Lower is Better
TX Costs	Logical Reads/TX	Lower is Better
Time	All Metrics	Lower is Better
Rates	Index Read Efficiency	Lower is Better
Rates	TX/Sec	Higher is Better
Rates	SQL/Sec	Higher is Better
Ratios	All	Higher is Better
Percentages	SQL Sync Read %	Higher is Better
Percentages	XML Sync Read %	Higher is Better
Percentages	Async Write %	Higher is Better
Percentages	Sort Ovflo %	Lower is Better
Percentages	% CPU Busy	Lower is Better

www.database-brothers.com
Phone: 512-249-2324

Step 10: Share your Success

To become a DBI Certified Performance Hero, send screen shots of performance trend graphs indicating improvements, or screen shots of grids indicating performance improvements, to hero@database-brothers.com along with your contact information and a short essay description of the steps you took and the results you achieved. Someone from DBI will contact you to review your achievements and make arrangements to send you your DBI Certified Performance Hero award kit.

And, of course, it goes without saying, but it would be prudent to repeat steps 1-9 repeatedly until no further tuning improvements can be obtained. Also be sure to explore all of the features and capabilities of Brother-Panther; only a fraction of its functionality has been covered by this introductory guide.

Summary

Brother-Panther is a robust product with many features and capabilities. Please read the product documentation completely. As you explore all of the capabilities of the product, you will discover many ways that Brother-Panther can help you administrate your databases and improve and sustain excellent database performance.

Here's a quick list of additional functionality and tuning suggestions:

- Review Partition performance in a multi-partition database to ensure the load is appropriately balanced across partitions
- Review Bufferpool Performance; tune to improve hit ratios and reduce physical I/Os
- Review Tablespace Performance; tune to improve read and write times and asynchronous pages read per request
- Review Table Performance; look for tables that need to be reorganized (high Overflows), and don't forget the catalog tables
- Review Table Performance; for tables having the highest Rows Written, carefully review index definitions to ensure there are no indexes with low cardinalities or skew
- Repeatedly review Statement Performance and work to reduce the relative execution costs of costly Statements
- Monitor Performance Trends over Time
- Explore the DBI Repository Database; it is well documented with many comments - you may find additional information there that is helpful to you