

## Leveraging DB2 Performance Expert for Business Intelligence

This presentation shows how IBM DB2 Performance Expert can be used to analyze and fix typical performance issues that occur in Business Intelligence environments.

In this presentation, we'll use a hypothetical company, JK Superstore, to illustrate how DB2 Performance Expert can help analyze and address typical performance challenges. JK Superstore is a growing national retail chain that uses Business intelligence to find ways to increase profits by offering appropriate promotions. They have a distributed data warehouse that contains the data they analyze with BI tools. They have a mixed environment with multiple applications. Because of the large amount of data residing on different DB2 servers, they use database partitioning (DPF). They have been experiencing some poor performance with some of the analytical queries being run.

There are several challenges associated with troubleshooting and fixing performance problems. The first is that data needs to be captured from different sources, such as the various DB2 databases and operating systems. This can be especially challenging in the distributed data partitioned environment where there is a lot to monitor and correlate. Another issue is that performance data is sometimes difficult to interpret, especially if changes need to be tracked over time. Troubleshooting can be challenging because once a performance problem occurs, the relevant information for troubleshooting may no longer be available. Finally, isolating a query that is causing a performance problem is difficult in the mixed, many application, production environment.

JK Superstore plans to use DB2 Performance Expert to help them address their performance challenges. DB2 Performance Expert provides a rich set of performance monitoring capabilities for the BI environment. DB2 PE is based on a 3 tier monitoring architecture. Tier 1 is the monitored DB2 system. Tier 2 is the DB2 PE server that remotely monitors tier 1. Tier 3 is the DB2 PE client that connects to the server and performs all the monitoring tasks.

DB2 Performance Expert has several features that help address the performance challenges listed earlier. For example, with DB2 PE there's a single point of control for monitoring and analyzing the performance of DB2 databases and operating systems. DB2 PE can monitor a multi-partitioned DB2 system, showing performance metrics for each partition or aggregated. DB2 PE provides at a glance visualizations in the form of dashboards and diagrams, tailored for BI and showing changes over time. DB2 Performance Expert continually monitors BI systems and preserves performance data for analysis and reporting. It has built-in exception processing which allows the DBA to work in a management by exception mode. Finally, DB2 PE also provides a rich set of SQL tracing and analysis features that allow identification of problem queries and of the applications that run them.

Now we'll illustrate a few of the ways DB2 PE can address some typical performance challenges. As stated earlier, some of the analytical queries run at JK Superstore are experiencing poor performance. To find out more about the poor query performance, the administrator can start by identifying the top long running queries. Using the dynamic SQL statements window, he can find the statements that take the longest execution time or are the top statements with regard to other critical indicators such as CPU time.

Using DB2 Performance Expert's Warehouse Query Facility function, it is possible to identify and then to drill down to the details on the applications running the queries with slow performance. If the application is no longer running, the administrator can see the relevant information in the history information that DB2 PE keeps.

At this point the administrator can check for skews. A skew is an uneven distribution of behavioral or residual distribution of performance indicators over the distributed system components. Skews are potential reasons for performance problems. For the particular application of interest, the administrator can see skews over the different partitions in key indicators such as user CPU time, system CPU time, or sort overflows. A skew might occur because there's more data on one partition than on the others. DB2 PE can be used to look at space management data for further analysis.

DB2 PE allows the administrator to see and get details on the SQL statement plan. By looking at the table queue component of the statement plan, the administrator can determine whether the model needs to be revised (by changing the partition keys or the distribution of tables to node groups) or whether the query needs to be revised. The administrator can also bring up the DB2 Index Advisor for assistance in determining how to revise the model.

After the analysis, the administrator may determine that the particular query is having a severe negative impact on the performance of the system and needs to be cancelled. The administrator can do that right from the DB2 Performance Expert client. The associated application will be rolled back and closed immediately.

DB2 Performance Expert's dashboards allow at-a-glance overall monitoring of system health. DB2 PE comes with several BI-tailored dashboards consisting of data views that track key BI performance indicators. Importing these data views results in a quick setup of the following dashboards: CPU/Memory/Disk, fast communications manager (FCM), Incremental Load, Page I/O, Sorts, Storage and Workload.

DB2 Performance Expert's dashboards allow at-a-glance overall monitoring of system health. DB2 PE comes with several BI-tailored dashboards consisting of data views that track key BI performance indicators. Importing these data views results in a quick setup of the following dashboards: CPU/Memory/Disk, fast communications manager (FCM), Incremental Load, Page I/O, Sorts, Storage and Workload.

Dashboards allow for online monitoring of key performance indicators as well as historical browsing. Each diagram (or data view) shows a key performance indicator. The administrator can get a better understanding of what's going on with the system by seeing these views at once. This BI recommended dashboard focuses on CPU, memory and disk usage. In this particular dashboard, CPU utilization shows a lot of idle time and low I/O wait time. Also the run and block queues are rather short. Overall, it shows there is little activity on the system at the moment.

Using a dashboard focused on Sorts, the administrator can monitor and detect skews in the sorts. In this dashboard, the rows show various sort views, and the columns show the data for each partition. Partition 1 thru 5 contain data and therefore should look identical if there are no skews. The pattern within each row looks the same except for a difference in sort heap for partition 5.

The administrator doesn't have to continually monitor the key performance parameters. Because DB2 Performance Expert tracks exceptions, the administrator can check for them on the system overview panel or request to be alerted (via email) when they are found. For example, the administrator can be alerted whenever TempSpace, a key parameter, fluctuates above 80%. In the situation shown above, high asynchronous reads (i.e., prefetch ) from the DB2 side correlate to high I/O on the operating system. An exception was reported for high table space usage; the administrator can drill down on it to get more details.

By using the BI templates supplied with DB2 Performance Expert, the administrator has a head start on setting appropriate thresholds for the key performance indicators.

This presentation has illustrated at a high level how DB2 Performance Expert can help with typical BI performance challenges and tasks. In particular, we touched upon understanding long running queries, detecting skews and dashboard monitoring of key performance indicators. Some of the other situations where DB2 Performance Expert can assist the BI administrator are the following: Checking to see if the system is CPU bound, ensuring that sorting is done in an optimal way, the ability to see and tune page I/O metrics, SQL tracing capability of an application or an entire BI system to understand performance problems. DB2 PE also provides the means to monitor and tune load processes, it provides parameter marker checking to optimize statements, helps verify materialized query table effectiveness and also allows for tuning of the FCM.

For more information on how DB2 Performance Expert can be used for Business Intelligence performance management, see the whitepaper on this topic, available on the IBM DB2 Performance Expert website:

<http://www.ibm.com/software/data/db2imstools/db2tools/db2pe/db2pe-mp.html>

It contains details on how to use DB2 Performance Expert in the various scenarios touched upon in this presentation

For further information on this tool and other tools from IBM, visit the DB2 and IMS Tools website: <http://www.ibm.com/software/data/db2imstools>

Thank you.