

# End-to-End Performance Monitoring of Databases in Distributed Environments

Stefanie Scherzinger, Holger Karn, Torsten Steinbach  
IBM Deutschland Research and Development, Böblingen  
{sscherz, holger.karn, torsten}@de.ibm.com

**Abstract:** This demonstration features the IBM *DB2 Performance Expert for Linux, Unix and Windows*, a high-end database monitoring tool that is capable of end-to-end monitoring in distributed environments. Performance Expert closely tracks the execution of database workload, from the point when the workload is issued in a database application, to its actual execution by the database server. This enables us to break down the response time as experienced by the user into time spent waiting for connections, within database drivers, the network, and finally, within the database server. From a data management point-of-view, the main challenge here is the handling of large amounts of performance data. We show how this challenge is met by continuous aggregation, and the definition of so-called *workload cluster groups*, which organize and narrow down the data of interest.

## 1 Motivation

Whether configuring, optimizing, or administrating databases and database applications, it is vital to maintain statistics on the database health and performance. The data collected during database monitoring is the basis for long-term planning of resources, as well as for short-term problem determination.

For instance, consider users who complain about long response times of a database application. In a typical customer environment, Java<sup>†</sup> applications access a database remotely. Locating the causes of performance problems in such a distributed setting can be a serious challenge. The response time, as experienced by the users, can actually be broken down into the time spent inside the application server, the database drivers, the network, the operating system, and finally, the database server. Within the individual layers, it can be helpful to further itemize the costs for particular tasks, such as locking or I/O.

Performance bottlenecks may lurk in any of these layers. While the database administrator (DBA) is often among the first in line for pinpointing a performance issue, the DBA typically only knows the database side of the problem.

There is a variety of dedicated tools for monitoring each of these layers, such as for monitoring the performance of the operating system, the network traffic, or the database itself. Yet when the layers are observed in isolation, a considerable amount of time is spent on consolidating the views of different monitoring tools.

For instance, when Java applications do not register with the database, DBAs may find




Figure 1: Breakdown of response times in a rainbow chart.

it hard to map these seemingly nameless applications to the applications that users are complaining about. However, if network problems are responsible for delays, monitoring the database in isolation will not reveal the problem. Another example concerns connection pooling in applications. If a connection pool is exhausted, users experience sluggish performance, yet the DBA cannot find anything wrong with the database.

What is missing here is a holistic view on the execution of database workload from the one end, that of the application, to the other, namely the database server. That is, we need an *end-to-end* approach to database performance monitoring.

## 2 End-to-end Performance Monitoring

In this demonstration, we present end-to-end performance monitoring for database applications. This approach is featured in the *Extended Insight Feature* of IBM's *DB2 Performance Expert for Linux, Unix and Windows* [IBM08, KBS08]. DB2 Performance Expert is a high-end monitoring tool for DB2\* installations. The initial release of the Extended Insight Feature focuses on WebSphere\* and other Java† applications accessing DB2 data on Linux†, Unix†, or Windows† platforms.

DB2 Performance Expert closely tracks the propagation of database workload through the major execution layers, so that DBAs can readily see where the processing time is spent. A graphical user interface features *rainbow charts*, as shown in Figure 1. The x-axis depicts the time passed during monitoring, whereas the y-axis shows runtimes in seconds.

The rainbow chart provides a comprehensive view by breaking down the database transaction response time into the time spent inside the application, wait times inside connection pools, the processing inside database drivers, the network, and finally, the time spent within the database server itself. This allows the administrator to observe the end-to-end runtime behavior of the incoming database workload over time.

\*Trademarks of IBM in USA and/or other countries.

†Other company, product, or service names may be trademarks or service marks of others.

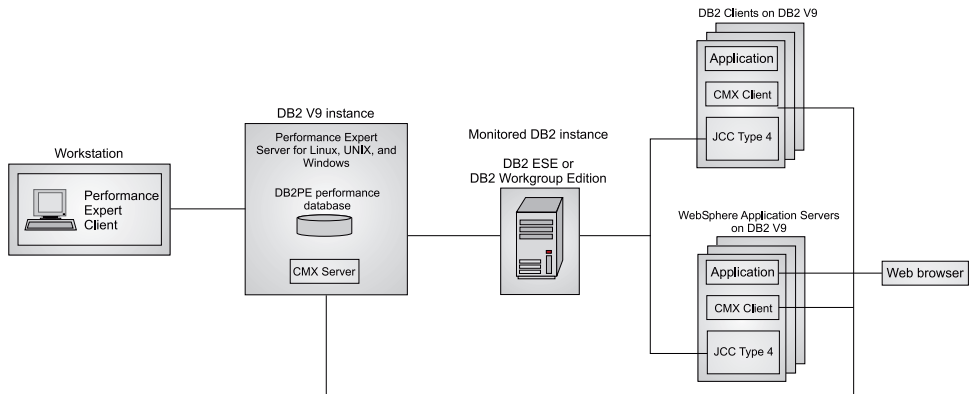


Figure 2: Distributed system architecture.

**Data management challenges.** The handling of large amounts of performance data with little overhead to the monitored system poses a challenge. End-to-end performance statistics are generated at the level of single transactions as well as the execution of single SQL statements. In DB2 Performance Expert, this data management challenge is effectively addressed in two ways.

- Users define *workload cluster groups* based on DB2 connection attributes. Performance data is then filtered and aggregated according to the DBA's interests. For instance, the data can be clustered by the names of users who issued the workload, and filtered for specific applications.
- Historical monitoring data is repeatedly aggregated over specific *data retention periods*. DBAs can access performance data months and even years later. Historical data is aggregated over longer timeframes, and thus more compactly, than more recent performance data.

**System architecture.** Figure 2 shows the distributed system architecture of DB2 Performance Expert. A DB2 instance is shown in the center. This instance is remotely monitored by an installation of *Performance Expert Server*. All monitoring data is stored persistently in the DB2PE performance database, maintained by *Performance Expert Server*.

DBAs interact with the *Performance Expert Client*, a graphical user interface that remotely accesses the performance database and visualizes its contents. By defining workload cluster groups, DBAs obtain customized views over the end-to-end monitoring data.

End-to-end monitoring is shown for two Java applications. These generate database workload on the monitored DB2 instance. The Java applications need not be recompiled or even relinked for end-to-end monitoring. Rather, they are merely configured to use JCC Type 4 as their database driver, and have a so-called *CMX Client* (short for *Client Management Extension*) installed on site.

The *CMX Clients* are responsible for tracking and collecting end-to-end monitoring data. They perform a pre-aggregation step on the collected data, and forward it to the *CMX Server* at regular intervals. The *CMX Server* is part of *Performance Expert*, and stores the end-to-end statistics in the performance database.

### 3 Demonstration Summary

In our demonstration, we set up an environment according to Figure 2, where we (1) generate database workload with a Java application, (2) monitor the execution of the workload on a DB2 instance, and (3) integrate these observations with end-to-end monitoring data. Then we show how users may systematically search out performance bottlenecks:

- We demonstrate how DBAs may define workload cluster groups. These are the means to view end-to-end monitoring data for specific applications, users, or clients.
- We show the breakdown of the response time in rainbow charts, as displayed in Figure 1. Rainbow charts are a comprehensive means for DBAs to quickly identify processing layers that may host performance bottlenecks.
- In case the database is harboring performance problems, DBAs can exploit the classic monitoring facilities of DB2 Performance Expert. These are custom-tailored towards performance tracking of DB2 installations.
- We query for the top- $k$  SQL statements, sorted by crucial criteria such as the end-to-end response time.
- We show how the characteristics of different data servers may be compared with regard to the average network time or the number of opened connections.
- We generate database workload for live-monitoring, but also show the stepwise aggregated historical performance data. History data is crucial for recognizing long-term changes in the system.

In summary, this demo will feature the monitoring functions of *DB2 Performance Expert for Linux, Unix and Windows*, with a focus on its novel end-to-end monitoring capabilities.

### References

- [IBM08] IBM. “DB2 Performance Expert for Linux, Unix and Windows”, 2008. <http://www-01.ibm.com/software/data/db2imstools/db2tools/db2pe/db2pe-mp.html>.
- [KBS08] Holger Karn, Ute Baumbach, and Torsten Steinbach. “Sneak peek: End-to-End Database Monitoring using IBM DB2 Performance Expert”. *IBM Database Magazine*, (4), 2008.

**Acknowledgements.** This work is the combined effort of our colleagues from the *DB2 Performance Expert for Linux, Unix and Windows* teams in Böblingen and Krakow. Many thanks.