

Revolution in IBM DB2 Performance: IBM DB2 Analytics Accelerator

by Namik Hrle, Ruiping Li, and Wolfgang Hengstler
IBM Boeblingen Development Laboratory and IBM Silicon Valley Lab

IBM DB2 Analytics Accelerator marks an inflection point in developing the DB2 technology. It brings lightning-fast performance to data-intensive and complex DB2 queries. This white paper covers the key DB2 Analytics Accelerator design and operational aspects that enable DB2 for IBM z/OS clients to benefit from faster performance, reduced CPU usage, and lower costs.

Abstract

DB2 Analytics Accelerator has been designed and developed as an internal DB2 component, so that in many aspects it is experienced as just another super-fast DB2 access path. The complex and data-intensive queries that characterize data warehouse, business intelligence, and analytics workloads can be now executed hundreds of times faster than before.

This paper answers several questions: What is business analytics, why is it important to your organization, and how does the IBM solution differ from others? Then, we address the unique DB2 Analytics Accelerator architecture that delivers a deep integration into DB2 and also takes advantage of the extremely fast, data-intensive query engine provided by IBM Netezza® technology. We also describe which query types qualify for the new “access path” and explain how to operate and monitor the accelerator.

Objectives

After defining business analytics and its importance to business needs, then addressing the key design and operational features, we will examine DB2 integration, when a query can be accelerated, and some performance monitoring features. Finally, we will discuss how DB2 Analytics Accelerator achieves high performance for the queries. Our goals are to

- Describe DB2 Analytics Accelerator architecture
- Learn how to use the accelerator, control acceleration, and maintain its content
- Learn how to interpret the new access path and monitor query acceleration

- Understand the query execution technology that powers the accelerator
- Learn which workloads and query types apply for the new access type

Introduction

Business analytics play a crucial role in today's workplace. The performance and cost of the DB2 Analytics Accelerator opens up unprecedented opportunities for enterprises to make use of the data on the IBM System z platform. Customers have seen dramatic improvement in the response times of the qualifying queries in some real, production-sized benchmarks. Running DB2 Analytics Accelerator on System z can result in some significant reductions in CPU usage. Of course, individual results will vary and depend on many other factors.

What Is Business Analytics? Timely, Accurate, and Secure Access to Business Information

Since the early days of data warehousing, the common statement from every vendor and pundit was that decision systems and transactional systems were vastly different and required separate platforms. Those days are over!

With the wealth of data available today, organizations are no longer willing to relegate information to the back office. Modern organizations are demanding access to customer purchase histories, customer behaviors, and trends of product sales at the time of contact—at the time of sale. This creates new challenges, because it is not enough for an enterprise to capture this data, but to process and transform these massive amounts of data into actionable knowledge. And this needs to be done quickly while the information is still relevant.

Data transformed into intelligence gives you more than a window into your current operations. It provides a likely view of the future—what is just around the corner and even further down the road. It helps leaders know with confidence all that has happened, is happening, and might happen to every aspect of the enterprise. Spotting the key patterns, extracting critical insights from data, and taking latency and cost out of making and implementing the right decision is what is defining industry leaders.

The world we live in today is increasingly instrumented, interconnected, and intelligent. We are experiencing a revolution, and information is at the heart of it. Businesses that are taking advantage of this new wealth of information are able to make more intelligent decisions and are rising to the top. They're managing large volumes of information in real time, incorporating analytics and predictive modeling, pervasively collecting and sharing information across the entire value chain, and speeding time to value by delivering trusted, accurate, and timely information to the right decision makers.

A company's survival can depend on the age of the data used to obtain an answer to critical business questions. With slow sales cycles, cutbacks, reluctant clients, and intense competition, business leaders are really feeling the heat to act and act fast, but a single bad decision today can be disastrous.

So what is the key to working smarter? It is having the right information and insight at the right time to drive smarter business outcomes. Working smarter means that your front-line business leaders receive timely information to uncover the new revenue opportunities and identify which product or service offerings are most likely to address the market requirement. It means business analysts can quickly access the right data points to evaluate key performance and revenue indicators in building successful corporate growth strategies. And, it means corporate risk and compliance units can recognize potential regulatory, reputational, and operational risks before they become realities.

The DB2 Analytics Accelerator gives your organization the speed to create the insights it needs to work smarter in this challenging environment. By putting the right answers in the hands of your decision makers immediately, DB2 Analytics Accelerator puts your business in the best position to quickly adapt and grow to answer the questions of tomorrow.

How Business Analytics Can Help Your Organization

Many organizations realize the benefit of improving business outcomes and improved decision making. The use of business intelligence and analytic applications is well understood to help make smarter decisions, achieve better results, and gain a deeper understanding of trends, opportunities, weaknesses, and threats. Organizations want to further analyze their data to gain additional insights into their business.

Today, however, the enterprise warehouse environment of an organization is facing many challenges. One such challenge is that the amount of data being stored in a typical warehouse environment is increasing. As the amount of data increases and sometimes the format of this data changes, the warehouse and end-user experience can be affected. It can become challenging for an organization to see the right information in an appropriate format and in the right timeframe for it to use in its analysis and decision-making process. Moving large amounts of data from disparate source systems to a warehouse can be a resource-intensive task. The increasing amount of data in some warehouses can also further affect any longer-running queries and reports that might exist in an organization. These slow-running queries, when executed with other mixed online transaction processing (OLTP) and online analytical processing (OLAP) workloads, can negatively affect the experience of existing users and cause further lack of acceptance for potential new users. Combine this with typical corporate priorities to become more productive, agile, and innovative, and it becomes more challenging to deliver on the promises of data warehousing and business analytics.

For many organizations, the concept that some of their longer-running DB2 for z/OS queries can be routed to an accelerator for processing is a plus. These queries may be in the form of batch SQL jobs or may be generated by means of corporate analytic and business intelligence (BI) tools—for example, ad hoc reporting from IBM Cognos® BI. The query accelerator available for DB2 for z/OS, which makes use of IBM Netezza technology, can make a big difference in the execution time of an analytic and warehouse

type of workload. Combining the benefits of both DB2 for z/OS (for OLTP-type queries) and DB2 Analytics Accelerator (for longer-running analysis queries) ensures that resources are shared appropriately for all warehouse users.

The DB2 Analytics Accelerator would likely benefit an organization that fit one of the following profiles:

- Wants to undertake a new reporting initiative on IBM System z to gain more insights
- Wants to consolidate disparate data to its existing System z platform while benefiting from integrated operational BI
- Wants to modernize an existing data warehouse and BI workload on System z

These types of organizations, with the appropriate workload, would likely see their elapsed time for longer-running queries being significantly reduced. They would also likely see their CPU usage on the mainframe being reduced, allowing DB2 for z/OS to focus on efficiently running their OLTP queries. Other benefits for these organization profiles are discussed in the following sections.

New System z BI initiative to gain more insight

This profile describes a System z organization that has identified a new reporting or operational BI initiative to analyze data that is not being currently analyzed. The organization would like to gain insights into the data and its business, while benefiting from having accelerated performance for complex analytics and queries. In this situation, it makes sense to use the DB2 Analytics Accelerator component for DB2 for z/OS. BI and analytic applications such as Cognos BI only need to connect to DB2 for z/OS and can still benefit from query acceleration.

The benefits of using DB2 Analytics Accelerator for a new reporting or operational BI initiative on System z include:

- Improved data insights for the organization's business users and business processes
- Performance, availability, and scalability benefits by blending System z and the DB2 Analytics Accelerator
- Acceleration benefits that are transparent to DB2 applications
- Simplicity and time to value for new mixed BI workload initiatives (OLTP and OLAP and analytics)

Consolidating disparate data to System z

This profile describes an organization that has created its data warehouse on System z and also has a number of disparate data marts (or islands of data) scattered around the organization, where some of its workload queries are executed. Some of these silos of information may be custom-built applications, which typically require ongoing maintenance and modification. There may be only a select few in the organization who are able to maintain or use some of these silos, and reporting might require some manual data

manipulation. The organization might have identified some potential benefits if some of the data flows and transformations to and from System z were eliminated, and if the organization wants a high performance integrated OLTP and BI analysis environment.

This type of organization could be facing any of the following challenges:

- There are multiple versions of the “truth.” This could include different applications providing different answers for the same information request, or different areas of the organization that own their own reporting data marts and apply their own interpretation of business rules.
- Corporate reporting and business analysis requires the use of multiple applications.
- Administration and management of multiple platforms and complex data integration processes are required.
- The value of consolidating data into a single easily managed platform (integrated OLTP and analysis/OLAP) has been identified, but some concerns may exist as to how analytic and traditional business intelligence workloads may perform on the mainframe.
- Deploying new data marts within the organization takes too long. Business benefit and value to the organization is not achieved in a timely manner.

The benefits of consolidating data on System z and including query acceleration with DB2 Analytics Accelerator are the same performance benefits mentioned in the previous organization profile. In addition, this type of organization might realize benefits including:

- Consolidated islands of data to a single secure data environment, providing “one version of the truth”
- An integrated OLTP and BI environment, enabling application queries that are required to use more real-time data
- Fewer servers to administer and less-competitive platforms
- The possible elimination of some network components, meaning fewer points of failure
- The enablement of data analytics consolidation through DB2 Analytics Accelerator
- The benefits of System z performance, scalability, and reliability combined with the accelerated performance of DB2 Analytics Accelerator
- The use of DB2 Analytics Accelerator to improve analysis workload performance, rather than requiring additional System z Integrated Information Processors to support the consolidated data warehouse environment

Modernizing an existing traditional BI workload

This profile describes an organization that has already created their data warehouse on System z. The warehouse contains historical data and coexists with many of the organization’s operational applications. The organization wishes to improve the performance of its existing BI and analytic workload.

Organizational challenges may include:

- Difficulty in extending the use of operational data for business analysis, embedding operational analytics in other applications, or daily business intelligence reporting.
- Long-running DB2 for z/OS queries. These queries may be executed from a BI environment and provide important business information. Currently, the queries can be scheduled in batch processes overnight so that they don't affect corporate users during the day. However, the overnight schedules could mean that information is not available in a timely manner—or that the full potential of having this information for other business processes is not realized.
- Forgotten queries, which due to performance issues are no longer executed. Some of these queries may have already been through exhaustive tuning efforts without success. If they were able to run successfully in a timely manner, the results could provide important decision-making information.
- Performance challenges with complex and ad hoc queries. Users, when building ad hoc queries through BI tools, may not realize the impact of their ad hoc querying.

The benefits of query acceleration using DB2 Analytics Accelerator for this organization include:

- Significantly improved query performance and execution time for individual queries or overall workloads, freeing up millions of instructions per second (MIPS) and storage space, therefore reducing processing cost
- The ability to execute queries that were either forgotten or blocked previously by the administrator due to performance issues
- Increased organizational agility by being able to more rapidly respond with immediate, accurate information and deliver new insights to business users
- Consolidated reporting on System z, where the majority of the data being analyzed lives, while retaining System z security and reliability

Impact on total cost of ownership (TCO)

In our scenario, query and reporting constitutes the DB2 dominant workload. In general, the DB2 Analytics Accelerator potential to effectively improve response times and possibly reduce costs by a CPU reduction is related to the costing model in effect in your organization. Most customers use monthly license charge (MLC) software based on a four-hour rolling peak average across a month. You must have a clear understanding of the way CPU is used and how CPU use for dynamic queries is reflected in your TCO.

Key Design and Operational Features

Figure 1.1 describes an IBM DB2 system, including applications, tools, and DB2 itself. Inside DB2 we have some familiar components, such as the Data Manager, Buffer Manager, Log Manager, Internal Release Lock Manager (IRLM), Relational Data System

(RDS), and more. Applications interact with DB2 through the application interfaces using SQL. Database administrators interact with DB2 through the operation interfaces, such as commands and utilities or performance monitoring and tuning tools.

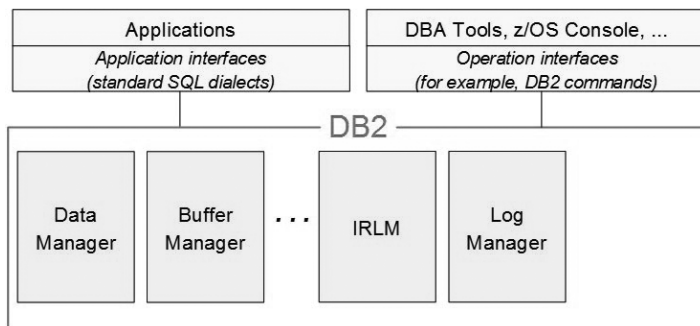


Figure 1.1: IBM DB2 components

Figure 1.2 illustrates that IBM now has a new “virtual” DB2 component, called IBM DB2 Analytics Accelerator. DB2 Analytics Accelerator has hardware and software components, based on IBM Netezza technology, used to accelerate complex queries that typically are seen in analytics applications.

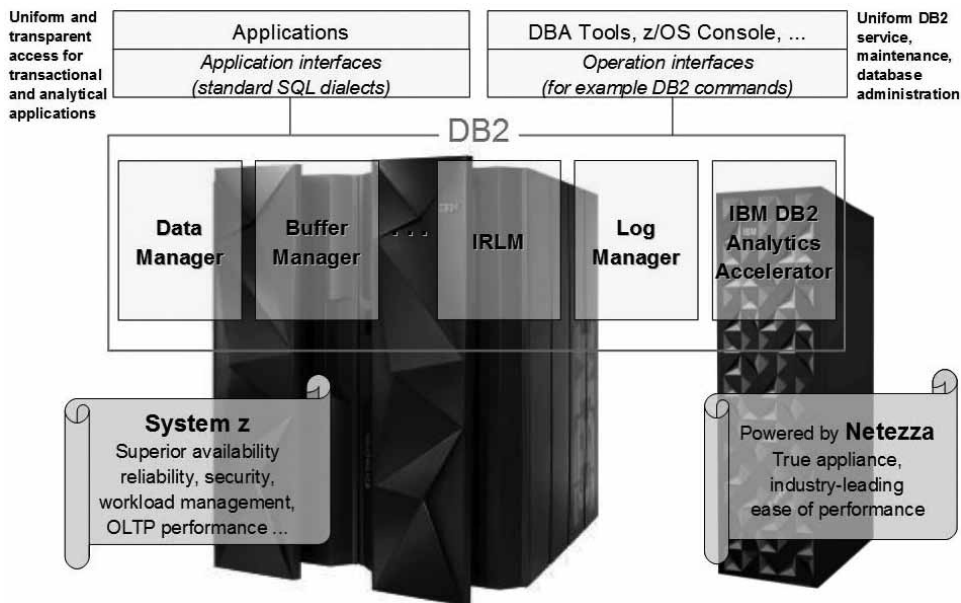


Figure 1.2: IBM DB2 becomes a Hybrid Database Management System

The IBM PureData™ System for Analytics appliance, connected to DB2, is enhanced to act as a DB2 accelerator. And DB2 has been enhanced with query acceleration to execute queries in DB2 Analytics Accelerator. DB2 Analytics Accelerator provides lightning-fast query performance transparently to the DB2 applications—at an affordable price. It opens up endless possibilities for new applications and workloads on data stored in DB2 for z/OS in enterprises.

Deep integration between DB2 Analytics Accelerator and DB2 for z/OS combines the best of both worlds into one single system. DB2 for z/OS is a world leader in OLTP, with superior availability, reliability, security, and serviceability. It also has world-class workload management capabilities. The PureData System for Analytics appliance provides superb data warehouse performance and the ease of use of an appliance.

Using it as an accelerator, it's not necessary to tend to administrative processes as with a standalone unit. You deal with data integrity and security on z/OS. DB2 Analytics Accelerator simply retains a copy of the data you want to accelerate the queries on and executes the queries for DB2.

DB2 Analytics Accelerator is administrated using a set of DB2 stored procedures. Query acceleration is viewed as a new query access path for DB2 which can be seen in the EXPLAIN output.

Query Execution Process Flow

Figure 1.3 illustrates a high-level query execution flow. There is an application on the left, DB2 in the middle, and DB2 Analytics Accelerator on the right. When the application submits a dynamic SQL query, DB2 will analyze it. If query acceleration is not enabled, or if the query does not qualify for acceleration, it will be executed locally within DB2.

If query acceleration is enabled and the query qualifies for acceleration, DB2 converts the query into Netezza syntax and routes it to DB2 Analytics Accelerator through an internal DB2 Analytics Accelerator to IBM DRDA® requestor interface. It talks to the DB2 Analytics Accelerator DRDA server on the SMP host in the PureData System for Analytics appliance, which completes the query execution within the appliance box and sends the result back to DB2 through DRDA. And the result is sent out to the application. The figure also shows the heartbeat messages from DB2 Analytics Accelerator to DB2, with DB2 Analytics Accelerator availability and performance indicators.

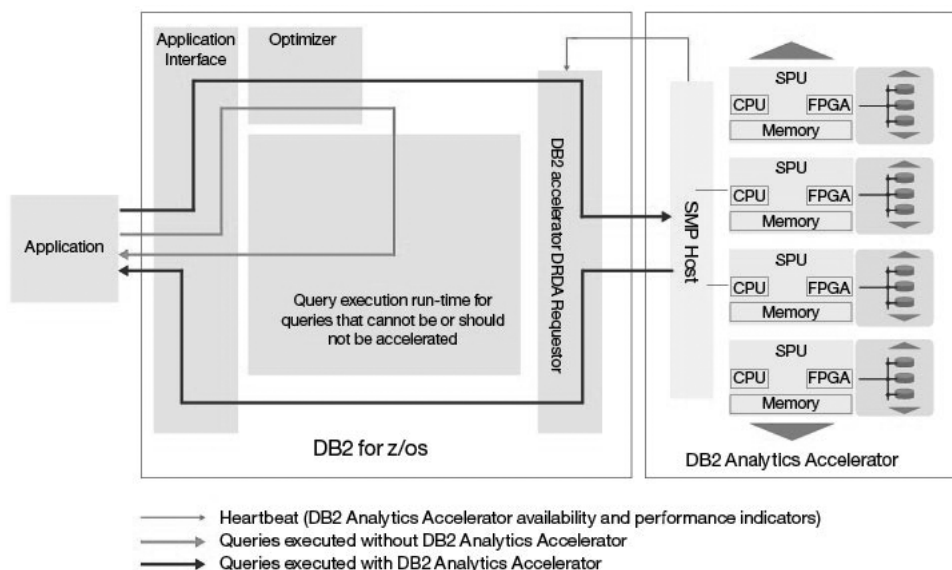


Figure 1.3: Query execution process flow

DB2 Analytics Accelerator Content Maintenance

Query acceleration by DB2 Analytics Accelerator is a new access path for DB2, just like an index access path. To enable this access path, you need to set up and enable query acceleration. One of the key steps is to have a copy of the table data in DB2 Analytics Accelerator so the queries can execute against them.

You define tables to be accelerated and then load the data from DB2 into DB2 Analytics Accelerator. You can refresh the data content periodically. The accelerator uses the DB2 UNLOAD utility to unload data, in parallel, to UNIX System Services (USS) pipes. DB2 Analytics Accelerator will read the data and convert it into a LOAD to the PureData System for Analytics appliance.

Partitions belonging to the same table can be loaded in parallel, with a user-controlled degree of parallelism to balance CPU and load throughput. Refresh or updates are done on a “per table,” “per partition,” or incremental basis.

Incremental update is a capability that enables tables on the DB2 Analytics Accelerator to be continually updated throughout the day. This technology reads the log of the database residing on DB2 for z/OS and applies those updates to the DB2 Analytics Accelerator. With this feature enabled, queries routed to the DB2 Analytics Accelerator will operate against a near real-time version of the data. It enables clients to dramatically lower the latency of data, letting decisions be made based on the most up-to-date information available. Customers use this feature when the workload being accelerated requires high currency of data for applications such as operational analytics. Incremental update is part of the accelerator’s integrated appliance form factor.