

IBM DB2 Analytics Accelerator for z/OS, Version 4.1

Delivering critical insight at the speed of business



The global economy has created a fiercely competitive environment where analytics-driven organizations considerably outperform their peers. In an effort to sustain that critical edge, companies constantly seek new ways to leverage the information available to them to identify opportunities, respond quickly to business situations and ultimately improve their competitive advantage.

Delivering the right information at the right time can help everyone make better business decisions—from executives to data analysts to front-line employees. Corporate agility and the ability to analyze information at the speed of thought are being viewed as core competencies and a source of competitive differentiation because the organization with better and faster insight is in the best position for success.

This need for deeper insight is generating new requirements for IT to:

- Provide operational decision makers with faster access to transactional data
- Improve query response times against larger volumes of data
- Improve the quality of the data provided to the business
- Control access to your most strategic asset: business data

To meet these requirements, many organizations are leveraging their existing IBM® zEnterprise® infrastructure to support expanding analytic needs. Business-critical analytic applications involve a broad mix of high-speed transactions and complex queries that must be processed simultaneously with very low latency. Organizations with zEnterprise have the distinct advantage of a hybrid solution that combines high-volume business transactions, batch reporting and complex analytic queries running concurrently in a mixed-workload environment.

The IBM DB2 Analytics Accelerator

The IBM DB2® Analytics Accelerator is a high-performance appliance that integrates zEnterprise and IBM Netezza® technologies to accelerate relevant data-intensive and complex queries. Complex multidimensional queries can run as much as 2,000 times faster than the same query running natively on IBM DB2 for z/OS®.

Together, DB2 Analytics Accelerator, DB2 for z/OS and zEnterprise form a self-managing hybrid environment that can run online transaction processing (OLTP), online transactional analytical processing (OLTAP) and online analytical processing (OLAP) workloads concurrently and efficiently. Organizations can obtain the results of analytic queries much more quickly than before, opening up unprecedented opportunities to use the data they already have on their zEnterprise platform. This hybrid infrastructure blends the best attributes of symmetric multiprocessing (SMP) using DB2 for z/OS with the best of massively parallel processing (MPP) using Netezza technology to deliver extraordinary performance at a reasonable cost.

Complex queries in seconds rather than hours

A large insurance company was running 300 mixed-workload queries against a table of over four billion rows. Of the 300 queries run against this table, 270 were completed in a second or sub-second timeframe. However, the company was challenged with the 30 queries that took minutes to hours to run. After deploying the DB2 Analytics Accelerator, the company was able to successfully accelerate the 30 “problem” queries; one ran 2,000 times faster without affecting the other 270.

The DB2 Analytics Accelerator provided a perfect blend of SMP and MPP processing to deliver the response the company needed. The insurance company’s short queries are processed by DB2 and the complex analytic queries are processed by the DB2 Analytics Accelerator, creating a solution that provides the best of both worlds.

Unprecedented response times for right-time analysis

When a query takes an unreasonable amount of time to provide the information necessary to make a key decision, the business opportunity may be lost. With the DB2 Analytics Accelerator, all of that changes. Queries that took hours can now complete in seconds. As part of its unique design, the DB2 Analytics Accelerator includes new breakthrough technologies to direct queries typically found in business intelligence (BI) and data warehousing applications, as well as operational applications that are infused with analytic intelligence, so that each query lands in its optimal setting for maximum speed and accuracy. It combines the zEnterprise qualities of service and Netezza hardware-accelerated technologies to speed up complex queries and deliver unprecedented response times in a highly secure and available environment (see Figure 1).

Organizations gain immediate value through this easy-to-deploy appliance. It transparently accelerates the most complex analyses from a wide variety of applications and tools such as IBM Cognos®, and it opens opportunities for new approaches and new applications. Organizations can now perform complex analysis in an environment based on industry-leading security and availability.

Application transparency

With DB2 Analytics Accelerator, there is virtually no need to change applications to achieve high-speed performance. In most cases, organizations just install the appliance, populate DB2 data into the Analytics Accelerator and begin reaping the benefits of high-performance decision making. Users and applications see only an IBM DB2 for z/OS interface; they can use their existing skills and applications without requiring the type of conversion or retraining normally associated with deploying new technologies.

The DB2 Analytics Accelerator requires no indices and no index maintenance. It accelerates workloads in a DB2 for z/OS environment with high-performance data warehouse query processing software and hardware that leverages massively parallel streaming hardware technology. It is optimized to provide fast response times for typical data warehouse or analytic-specific workloads, such as aggregations or large data scans.

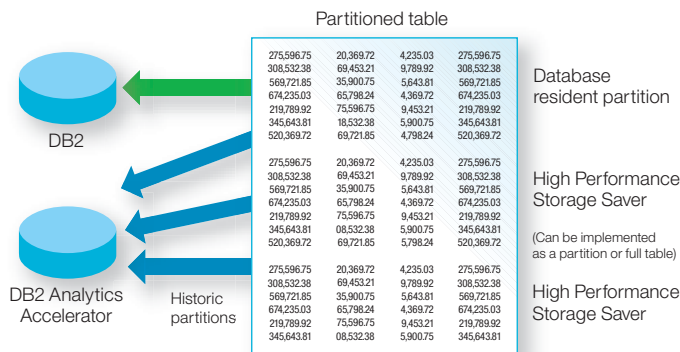
Choice of historical data location with High Performance Storage Saver feature

Most analytical systems are based on data that is 95 percent historical and therefore static. A retailer, for instance, might maintain seven years of past sales histories that contain every transaction for every product sold to every customer. Because this data is historical, it generally is not subject to revision or updates.

The High Performance Storage Saver feature of DB2 Analytics Accelerator helps reduce the cost of storing, managing and processing this type of data. Organizations can move static data within tables or table partitions to the DB2 Analytics Accelerator and remove the data from DB2 for z/OS (see Figure 2). All of the data is still managed by and safeguarded in the DB2 directory, and all of the queries that target that data are now directed only to the DB2 Analytics Accelerator. This process can dramatically reduce storage costs on zEnterprise and enable organizations to substantially increase the amount of history maintained for each subject area.



Figure 1. Blending IBM zEnterprise with IBM Netezza technologies.



- Historic partitions are stored only on DB2 Analytics Accelerator
- When partitions no longer require updating, you can reclaim the DB2 storage
- Place queries based on need: mixed-workload queries go to hotter-temperature partitions, aggregate-type querying goes to warm-temperature data

Figure 2. Each table or partition can be configured as DB2, DB2 Analytics Accelerator or both.

Real-time analytics with incremental update capabilities

The incremental update capability in the DB2 Analytics Accelerator enables administrators to continuously synchronize tables on the DB2 Analytics Accelerator with the data in the DB2 subsystem throughout the day. Part of the integrated appliance, this technology reads the log of the database residing on DB2 for z/OS and applies those updates to the DB2 Analytics Accelerator as they occur. Queries offloaded to the DB2 Analytics Accelerator will operate against a near-real-time version of the data, enabling dramatically lower data latency.

Access to current data allows decisions based on the most up-to-date information available—no need to use last week’s data to inform today’s strategies. Customers use this feature when the workload being accelerated requires extremely timely data for applications such as operational analytics.

High availability and workload balancing features

The IBM zEnterprise environment provides exceptional qualities of service and availability. DB2 Analytics Accelerator delivers high availability through an active/active architecture that also lends itself to significant user scalability.

DB2 Analytics Accelerator utilizes workload balancing to provide the most efficient use of all eligible Accelerators. If multiple DB2 Analytics Accelerators are connected to a DB2 for z/OS subsystem, DB2 checks the utilization of each Accelerator and then routes the query to the one with the lowest utilization. This capability also provides automated protection to help ensure system availability during both planned and unplanned outages. These features extend zEnterprise availability and workload balancing (based on IBM Parallel Sysplex® and DB2 data sharing architectures) to the DB2 Analytics Accelerator environment.

Direct interface with DB2 for z/OS

The DB2 Analytics Accelerator plugs into your DB2 for z/OS environment, complementing its traditional query processing. This direct interface enables deeply embedded capabilities behind the application layer and provides full transparency to applications that submit queries to DB2 for z/OS.

Data that is maintained on the DB2 Analytics Accelerator inherits all DB2 for z/OS data attributes, including security and recoverability. All data is loaded, backed up and retrieved through DB2 with no external connections to the DB2 Analytics Accelerator to minimize any opportunity for intrusion.

The DB2 Analytics Accelerator can speed up relevant dynamic and static queries, providing dramatic improvements in cost and performance, all transparent to the end user or applications. Because of this transparency, the DB2 Analytics Accelerator is fast and easy to deploy, and can be activated and deactivated by setting or clearing a software switch.

Query tuning elimination

Most organizations struggle with complex, long-running queries that become the bane of database administrators. They have invested hours, even days tuning and adding additional indices or materialized query tables (MQTs), only to decide that the query cannot be run within reasonable resources and must be

removed from the system. The DB2 Analytics Accelerator enables processing of these complex, long-running queries without requiring additional tuning or zEnterprise resources, such as indices or MQTs—increasing time-to-value while decreasing resource requirements.

IBM DB2 Analytics Accelerator: Technical highlights

Extreme performance for complex business analysis

- Exploits hardware accelerators (a multiengine field-programmable gate array [FPGA]) to deliver unprecedented query speed
- Marries the best of zEnterprise with Netezza technologies
- Automatically performs table scan reduction on compressed data without need for indices or compression directives

Database performance appliance

- Enables faster time-to-value through an easy-to-deploy, scalable appliance format
- Delivers unprecedented response times to enable “train of thought” analyses and analytics-enriched operational transactions
- Plugs into an existing DB2 for z/OS environment to dramatically improve query processing
- Provides transparency into applications that submit queries to DB2 for z/OS, without requiring application changes

Proven operational characteristics

- Extends zEnterprise manageability, security and availability to BI and data warehouse workloads
- Uses DB2 for z/OS to maintain data, taking advantage of its industry-leading capabilities in continuous availability and scalability through IBM Parallel Sysplex and data-sharing technologies

IBM DB2 Analytics Accelerator systems and sizes

N1001	002	005	010	015	020	030	040	060	080	100
Cabinets	¼	½	1	1½	2	3	4	6	8	10
S-Blades	4	7	14	18	28	42	56	84	112	140
Processing units	32	56	112	144	224	336	448	672	896	1120
Capacity (TB)*	8	16	32	48	64	96	128	192	256	320
Effective capacity (TB)**	32	64	128	192	256	384	512	768	1024	1280

N2001	005	010	020	040
Cabinets	½	1	2	4
S-Blades	4	7	14	28
Processing units	64	112	224	448
Capacity (TB)*	24	48	96	192
Effective capacity (TB)**	96	192	384	768

N2002	002
Cabinets	¼
S-Blades	2
Processing units	32
Capacity (TB)*	8
Effective capacity (TB)**	32

*Capacity: User data space

**Effective capacity: User data space with compression (4x compression assumed)

DB2 Analytics Accelerator: An integrated part of IBM zEnterprise Analytics System 9700 and 9710

DB2 Analytics Accelerator is part of the end-to-end IBM zEnterprise analytic solution. This comprehensive portfolio offers data warehousing, data transformation and integration, business intelligence, enterprise planning and predictive analytics, including DB2 for z/OS, InfoSphere® Information Server, InfoSphere Data Replication, Cognos BI for zEnterprise, DB2 Query Management Facility™ (QMF™), Cognos TM1, IBM SPSS® and zEnterprise Analytics System 9700/9710.

The zEnterprise Analytics System 9700 is a cost-competitive deployment model that integrates hardware, software and services to deliver business-critical analytics. This pretested, preselected and solution-priced offering includes DB2 for z/OS, DB2 Analytics Accelerator and zEnterprise hardware at its core. The offering also includes three optional add-on packs (Data Analytics Pack, Data Integration Pack and the Fast Start Services Pack) to provide enterprises with the flexibility to address their specific analytic needs.

For customers who require a smaller footprint solution, the zEnterprise Analytics System 9710 provides an entry-level, integrated offering built on the zEnterprise Business Class footprint to cost-competitively, quickly and securely deliver business reporting and analytics. The zEnterprise Analytics System 9710 offers the DB2 Analytics Accelerator and the three add-on packs as options.

These solutions give organizations that have data on zEnterprise the opportunity to further use the platform to support a complete and cost-effective analytics solution.

A deeper look inside

The DB2 Analytics Accelerator is optimized for data warehouse-type querying. Its efficient, hardware-accelerated query engine is designed to stream data to achieve maximum speed.

The DB2 Analytics Accelerator uses copies of the DB2 for z/OS database tables that are stored in the DB2 Analytics Accelerator. This combination leverages legendary DB2 for z/OS performance for transactional queries and industry-leading Netezza performance for analytical queries.

Analytical queries are performed by the IBM Netezza Asymmetric Massively Parallel Processing (AMPP) engine, an architecture that uses special-purpose hardware accelerators to decompress and filter data for relevance to the query before it is loaded into memory and given to the processor for aggregation. The result: elimination of the I/O wait time and processing overhead that normally result in query bottlenecks. This new approach also removes the need for tuning, and helps reduce the time required to design and maintain the indexes and materialized views, and to collect statistics whenever the data changes.

Breakthrough technology: Hardware acceleration

Each node (blade) in the DB2 Analytics Accelerator is an independent server containing multi-core CPUs. Each CPU core is teamed with multiple spindles of its dedicated storage drives, gigabytes of random access memory and a hardware accelerator with a multiengine field-programmable gate array (FPGA). These FPGAs are used for data decompression, data filtering and early SQL projections and restrictions. The nodes are configured in a shared-nothing environment, which means each of the CPUs has its own memory, and has no need to communicate with other nodes for record locking and lock resolution.

FPGAs accelerate processing of data before it reaches the CPU on each of the nodes (see Figure 3). Each FPGA completes its work with enormous efficiency, drawing little power and generating little heat. The focus of the system's optimization is to enable streaming processing—in effect, processing and retiring analytic operations as rapidly as the relevant information for them can be read from the many parallel disk drives in the system. The typical impact of this work is a reduction of 95 percent or more in the data required for further processing by the onboard CPU cores and memory.

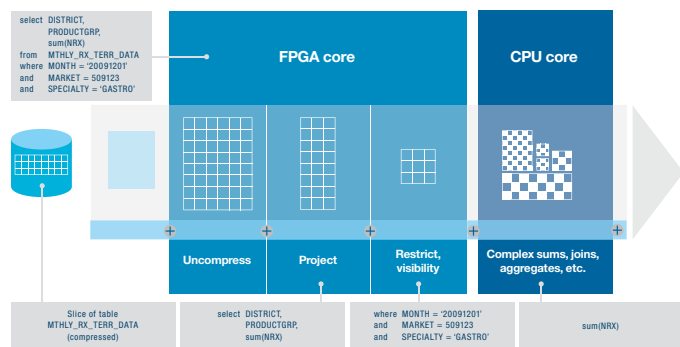


Figure 3. Each blade in the DB2 Analytics Accelerator has a CPU core teamed with an FPGA core with two to three active disks to accelerate processing of data before it reaches the CPU.

Designed to meet the needs of business and IT leaders in any industry

The DB2 Analytics Accelerator leverages the power of zEnterprise and DB2 for z/OS to offer significant value for any organization seeking to drive smarter business outcomes faster.

From dashboards to trend-based analytics, it provides high-performance analytics to support business users who need fast, accurate answers to critical business questions. The DB2 Analytics Accelerator delivers the answers you need, where and when you need them.

As a hybrid system, DB2 for z/OS with the Analytics Accelerator allows organizations to go beyond the traditional data warehouse or decision-support system by enriching their mission-critical operational applications with prescriptive analytics at the point of interaction.

Make smarter business decisions with flexible data warehousing capabilities

Maintaining a competitive edge means using all the information at your disposal to create actionable insights that drive smarter decision making and smarter transactions. Because data warehousing technology forms the foundation for business analytics systems, IT departments must ensure that their strategy and infrastructure align with their goals, as well as with the goals of the business as a whole. IBM data warehousing and business analytics solutions on zEnterprise provide an end-to-end solution on a single platform—from data warehousing to business intelligence to predictive analytics—that is capable of scaling to meet the breadth of business user requirements for complete and accurate business information. And it does this quickly and securely with outstanding availability and performance.

For more information

To learn more about the IBM DB2 Analytics Accelerator, please contact your IBM representative or IBM Business Partner, or visit: ibm.com/software/data/db2/zos/analytics-accelerator

To learn more about the IBM zEnterprise Analytics System models and how they work with the DB2 Analytics Accelerator, visit:

- ibm.com/software/products/us/en/zenterprise-smartana9700
- ibm.com/software/products/us/en/zenterprise-smartana9710

Additionally, IBM Global Financing can help you acquire the software capabilities that your business needs in the most cost-effective and strategic way possible. We'll partner with credit-qualified clients to customize a financing solution to suit your business and development goals, enable effective cash management, and improve your total cost of ownership. Fund your critical IT investment and propel your business forward with IBM Global Financing. For more information, visit: ibm.com/financing



© Copyright IBM Corporation 2013

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
October 2013

IBM, the IBM logo, ibm.com, Cognos, DB2, InfoSphere, Parallel Sysplex, QMF, Query Management Facility, SPSS, zEnterprise, and z/OS are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

Netezza is a trademark or registered trademark of IBM International Group B.V., an IBM Company.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.



Please Recycle