IBM DB2 Analytics Accelerator: A revolution in performance

By Namik Hrle, Ruiping Li and Wolfgang Hengstler

Contents

Executive overview

IBM® DB2® Analytics Accelerator marks a transitional point in developing DB2 technology. The complex and data-intensive queries that characterize data warehouse, business intelligence, and analytics workloads can now be run hundreds of times faster than before. DB2 Analytics Accelerator features key design and operational aspects that enable DB2 for z/OS® clients to benefit from faster performance, reduced processor use, and lower costs.

The introduction of DB2 Analytics Accelerator brought a revolutionary change:

IBM DB2 Analytics Accelerator for z/OS complements DB2 for z/OS, which is built for transactional workloads with a cost-effective high speed query engine to run complex business analytics workload. As part of its unique design, the DB2 Analytics Accelerator includes new breakthrough technologies to reroute queries typically found in transactional workloads to DB2 for z/OS and queries typically found in business intelligence (BI) and data warehousing applications to the integrated IBM PureData™ System for Analytics appliance. Thereby, DB2 Analytics Accelerator turns DB2 for z/OS into a universal database management system, capable of handling both: transactional and analytical data. DB2 Analytics Accelerator, together with DB2 for z/OS form a new, self-managing, hybrid database management system that runs every query workload in the most efficient way, so that each query lands in its optimal setting for maximum speed, execution and cost efficiency.

To understand the full value of DB2 Analytics Accelerator, a basic understanding must be reached regarding three foundational questions:

- What is business analytics?
- Why is it important?
- How does this solution differ from others?

Only by answering these questions can users begin to grasp the importance of the unique DB2 Analytics Accelerator architecture. With deep integration into DB2 and the extremely fast data-intensive query engine that is provided by IBM PureData for Analytics, based on Netezza technology,® DB2 Analytics Accelerator has the potential to transform how organizations interact with their DB2 systems.

Users can only begin to understand this revolutionary technology with a clear definition of business analytics, its importance to business needs, and an understanding of the key design and operational features of DB2 Analytics Accelerator. For those people that want even more information, they can explore more complex topics like when to accelerate queries and how to operate performance monitoring features. This additional information can help users understand exactly how DB2 Analytics Accelerator achieves its greater levels of performance.

The more users understand this tool, the more likely they are to gain several critical new skills:

- A holistic view of DB2 Analytics Accelerator
- The ability to use the accelerator, control acceleration, and maintain its content
- The capacity to interpret the new access path and monitor query acceleration
- An understanding of the query execution technology that powers the accelerator
- The knowledge of which workloads and query types apply to the new access path

Introduction

Business analytics plays a crucial role in today's workplace. The performance and cost of DB2 Analytics Accelerator create unprecedented opportunities for enterprises to use their data on the IBM System z® platform. Clients have seen dramatic improvements in the response times of qualifying queries in real, production-sized benchmarks. Furthermore, running DB2 Analytics Accelerator on System z can result in significant reductions of processor use.

Business analytics: Timely, accurate, and secure access to business information

Since the early days of data warehousing, the common sentiment from vendors and experts 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 systems 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. This increased demand for information creates new challenges. To meet these demands, an enterprise must not only capture these massive amounts of information, but also process and transform it into actionable insights. This processing must be done quickly while the information is still relevant.

Data that is transformed into intelligence offers more than a window into current operations. It provides a likely view of the future. It helps leaders know with confidence all that happened, is happening, and might happen to every aspect of the enterprise. Identifying key patterns, extracting critical insights from data, and taking latency and cost out of making and implementing the right decision are what define industry leaders today.

Today's world is becoming increasingly instrumented, interconnected, and intelligent. Companies are experiencing a revolution, and information is at the heart of it. Businesses that are taking advantage of this new wealth of information to make more intelligent decisions are rising to the top. They are managing large volumes of information in real time and incorporating analytics and predictive modeling. They are 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 that is used to obtain answers to critical business questions. With slow sales cycles, cutbacks, reluctant customers, and intense competition, business leaders are being pressured to act fast, but a single bad decision 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 front-line business leaders receive timely information to uncover new revenue opportunities and understand which product or service offerings are most likely to address the market requirement. It means that 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 that corporate risk and compliance units can recognize potential regulatory, reputational, and operational risks before they become realities.

DB2 Analytics Accelerator gives organizations the speed to create the insights they need to work smarter in this challenging environment. By putting the right answers immediately in the hands of decision makers, businesses can be in the best position to quickly adapt and grow toward answering the questions of tomorrow.

How business analytics can help your organization

Many organizations are aware of the benefits of improving decision making to elicit better business outcomes. The use of business intelligence and analytics applications is understood to help in making smarter decisions, achieving better results, and gaining a deeper understanding of trends, opportunities, weaknesses, and threats. Understandably, organizations want to further analyze their data to gain more insights into their business operations.

Today, however, the enterprise warehouse environment of an organization faces many challenges. One such challenge is that the amount of data that is being stored in a typical warehouse environment is increasing. As the amount of data increases and the format of this data changes, the warehouse and user experience can be affected. An organization can find it challenging to see the right information in an appropriate format and in the right timeframe for them to use it in their analysis and decision-making processes.

Moving large amounts of data from disparate source systems to a data warehouse can be a resource-intensive task. The increasing amount of data in some warehouses can also further affect any long-running queries and reports that exist in an organization. These long-running queries, when run with other mixed online transaction processing (OLTP) and online analytical processing (OLAP) workloads, can negatively affect the experience of existing users and cause further difficulties for potential new users. Combine this situation with typical corporate priorities to become more productive, agile, and innovative, and delivering on the promises of data warehousing and business analytics becomes much more challenging.

For many organizations, the concept that some of their long-running DB2 for z/OS queries can be routed to an accelerator for processing is a clear benefit. These queries might be in the form of batch SQL jobs, or they might be generated by corporate analytic and business intelligence (BI) tools, such as ad hoc reporting from IBM Cognos® Business Intelligence. The query accelerator that is available for DB2 for z/OS, which uses IBM Netezza technology, can make a large difference in the run time of an analytics or warehouse type of workload. Combining the benefits of DB2 for z/OS for OLTP-type queries and DB2 Analytics Accelerator for long-running analysis queries ensures that resources are shared appropriately for all data warehouse users.

DB2 Analytics Accelerator is likely to benefit organizations that fit one of the following profiles:

- Companies that are looking to undertake a new reporting initiative on System z to gain more insights
- Companies that are consolidating disparate data to an existing System z platform while benefiting from integrated operational BI
- Companies that are planning to modernize an existing data warehouse and BI workload on System z

Organizations within these profiles that have the appropriate workload are likely to see their elapsed time significantly reduced for long-running queries. They are also likely to see reduced processor use on the mainframe, allowing DB2 for z/OS to dedicatedly and efficiently run their OLTP queries.

New System z BI initiative to gain more insight

One profile fits an organization that is running on the System z platform and has identified a new reporting or operational BI initiative to analyze data that is not being currently analyzed. This type of organization wants to gain insights into their data and business processes, while benefiting from having accelerated performance for complex analytics and queries. In this situation, the DB2 Analytics Accelerator component for DB2 for z/OS makes sense. BI and analytics applications, such as Cognos Business Intelligence, need only to connect to DB2 for z/OS, and can still benefit from query acceleration.

Many benefits accompany the DB2 Analytics Accelerator for a new reporting or operational BI initiative on System z:

- Improved data insights for the organization's business users and business processes
- Performance, availability and scalability benefits from integrating System z and DB2 Analytics Accelerator
- Acceleration benefits that are not exclusive to DB2 applications
- Simplicity and time to value for new mixed BI workload initiatives like OLTP, OLAP, and analytics

Consolidating disparate data to System z

Another profile describes organizations that have created their data warehouses on System z and have several disparate data marts that are scattered throughout their systems. Some of these silos of information might be custom-built applications, which typically require ongoing maintenance and modification. This type of organization might have only a select few applications that can maintain or use some of these silos, and reporting might require some manual data manipulation. The organization might have identified potential benefits if some of the System z data flows and transformations were eliminated and if the organization wants a high performance-integrated OLTP and BI analysis environment.

This type of organization might face many challenges:

- A lack of consistency that is caused by applications that provide different answers for the same information request. This inconsistency could also be caused by various areas of an organization that might have their own reporting data marts and apply their own interpretation of business rules.
- Multiple applications for corporate reporting and business analysis
- Administration and management that are required for multiple platforms and complex data integration processes
- Concerns as to how analytics and traditional business intelligence workloads might perform on the mainframe despite understanding the value of consolidating data into a single platform
- Business benefits and value to the organization are not achieved in a timely manner because it takes too long to deploy new data marts within the organization

Consolidating data on System z and including query acceleration with DB2 Analytics Accelerator have the same performance benefits as undertaking a new BI initiative to gain more insight. In addition, this type of organization might realize its own distinct benefits:

- Consolidation of disparate data marts into a single secure data environment, providing a greater degree of consistency
- An integrated OLTP and BI environment, enabling application queries that require the use of real-time data
- Fewer servers to administer and fewer competing platforms
- The possible elimination of some costly, failure-prone network components
- 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 extra System z Integrated Information Processors to support the consolidated data warehouse environment

Modernizing a traditional BI workload

Some organizations have already created their data warehouses on System z. Their warehouses contain historical data and coexist with many other operational applications. This type of organization wants to improve the performance of its existing BI and analytics workload.

These particular organizations can face many critical challenges:

- 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 can be run from a business intelligence environment and provide important business information. The queries might be scheduled in batch processes overnight so that they do not affect corporate users during the day. However, the overnight schedules can 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 that are no longer run because of performance issues. Some of these queries might have already been through exhaustive tuning efforts without success.
- Performance challenges with complex and ad hoc queries. When building ad hoc queries through BI tools, users might not realize the impact of their ad hoc querying.

Query acceleration with DB2 Analytics Accelerator can provide these organizations with many benefits:

- Query performance and run time of individual queries or overall workloads can improve significantly, freeing up storage space and allowing for millions of instructions per second, ultimately reducing processing costs
- The ability to run queries that were forgotten or blocked previously by the administrator because of performance issues
- An increase in organization agility from the ability to more rapidly respond with accurate information and quickly deliver new insights to business users
- Consolidated reporting on System z, where most of the data that is being analyzed exists, while retaining System z security and reliability

Impact on total cost of ownership

Query and reporting can constitute the DB2 dominant workload. The potential of DB2 Analytics Accelerator to effectively improve response times and lower costs through processor usage reduction is directly related to an organization's specific costing model. Most clients use monthly license charge (MLC) software that is based on a four-hour rolling peak average across a month. Companies must have a clear understanding of the way a processor is used and how processor use for dynamic queries is reflected in their total cost of ownership (TCO).

Key design and operational features

Figure 1 describes an IBM DB2 system, including the applications and tools. Some familiar components, including Data Manager, Buffer Manager, Log Manager, Internal Release Lock Manager (IRLM), and RDS, are included in DB2.

IBM DB2 components

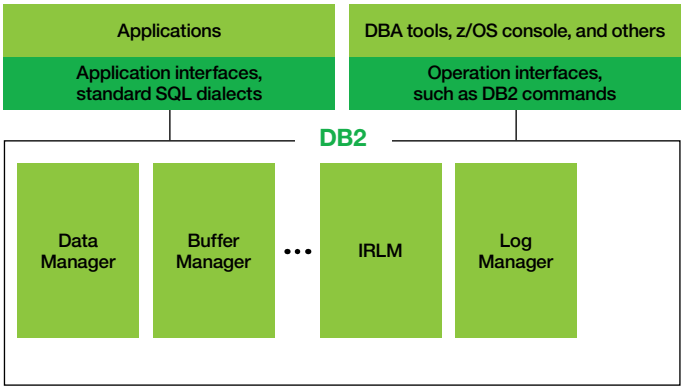

Figure 1: DB2 system with applications, tools, and components

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

Figure 2 illustrates that IBM now has a new virtual DB2 component, called IBM DB2 Analytics Accelerator. DB2 Analytics Accelerator has hardware and software components that are based on IBM Netezza technology and used to accelerate complex queries that are typically seen in analytics applications.

DB2 for z/Os approach: Hybrid database management system

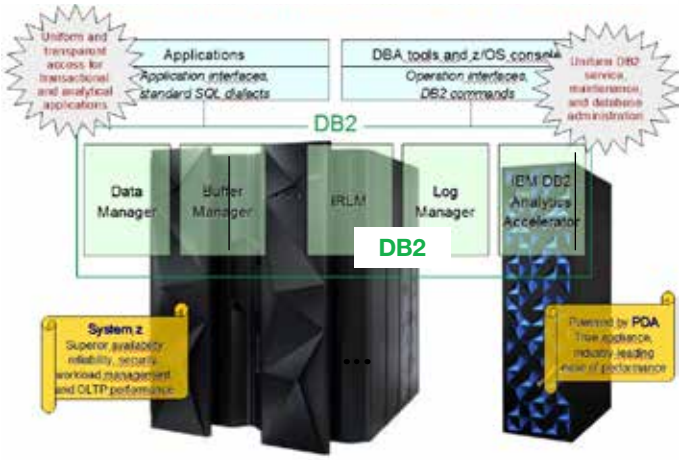

Figure 2: DB2 system with the DB2 Analytics Accelerator

The IBM PureData™ System for Analytics appliance, which is connected to DB2, is enhanced to act as a DB2 accelerator. With this solution, DB2 is enhanced with query acceleration to run queries in DB2 Analytics Accelerator. This solution provides fast query performance transparently to the DB2 applications at an affordable price. It also opens up endless possibilities for new applications and workloads on data that is stored in DB2 for z/OS.

With deep integration between DB2 Analytics Accelerator and DB2 for z/OS, this solution 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, in addition to world-class workload management capabilities.

When using the PureData System for Analytics as an accelerator, users do not need to tend to administrative processes as with a stand-alone unit. They can deal with data integrity and security directly on z/OS. DB2 Analytics Accelerator simply retains a copy of the data to accelerate the queries on DB2.

DB2 Analytics Accelerator is administered by using a set of DB2 stored procedures. Query acceleration is viewed as a new query access path for DB2, as can be seen with the

EXPLAIN output within DB2.

DB2 Analytics Accelerator query process and features

Figure 3 illustrates a high-level query process flow. When the application submits a dynamic SQL query, DB2 analyzes it. If query acceleration is not enabled or the query does not qualify for acceleration, then it is run locally within DB2.

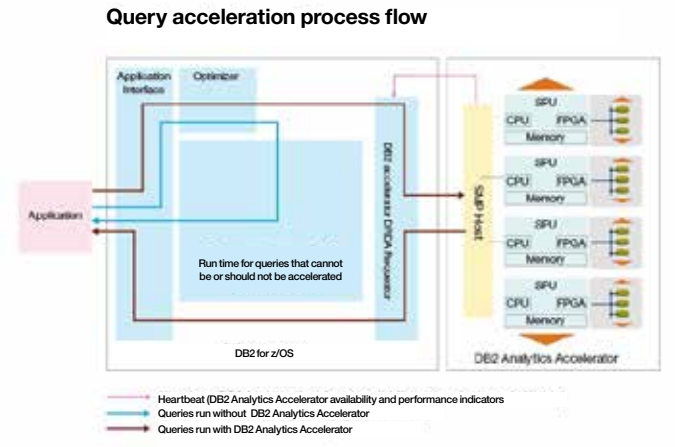

Figure 3: High-level query process flow that shows heartbeat messages

If query acceleration is enabled and the query qualifies for acceleration, DB2 converts the query into Netezza syntax. DB2 then routes it to DB2 Analytics Accelerator through an internal DB2 accelerator IBM DRDA® requester interface. DB2 communicates with the DB2 accelerator DRDA server on the symmetric multiprocessor (SMP) host in the PureData System for Analytics appliance. This SMP host completes the query process within the PureData System for Analytics appliance and sends the result back to DB2 through the DRDA path. The result is then sent to the application.

Figure 3 also shows the heartbeat messages from DB2 Analytics Accelerator to DB2 with DB2 Analytics Accelerator availability and performance indicators.

DB2 Analytics Accelerator content maintenance

Query acceleration through DB2 Analytics Accelerator is a new access path for DB2, just like an index access path. To enable this access path, users must set up and enable query acceleration. A key step in this process is to have a copy of the table data in DB2 Analytics Accelerator so that the queries can run against them.

Users define tables to be accelerated, and then load the data from DB2 into DB2 Analytics Accelerator. They can refresh the data content periodically. DB2 Analytics Accelerator uses the DB2 UNLOAD utility to unload data in parallel to USS pipes. DB2 Analytics Accelerator reads the data and then converts it into LOAD for the PureData System for Analytics appliance.

Partitions that belong to the same table can be loaded in parallel, with a user-controlled degree of parallelism to balance processor use and load throughput. Updates are run on a per-table, per-partition, or incremental basis.

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

DB2 Analytics Accelerator table definition and deployment

Before loading data, users must define the tables to be accelerated. All administrative tasks are achieved by using DB2 stored procedures for DB2 Analytics Accelerator. IBM DB2 Analytics Accelerator Studio provides a graphical user interface (GUI) to these stored procedures for convenience in performing the administrative tasks. Applications can start the stored procedures directly.

The stored procedures update the catalog tables, which provide the necessary information to support the query acceleration. To define and deploy tables into DB2 Analytics Accelerator, users identify the tables for which queries must be accelerated, and then load and enable the data for query acceleration.

High-performance storage saver

Most analytical systems are based on data that is over 95 percent historical and static. For example, a retailer might maintain seven years of past sales histories that contain every transaction for every product that was sold to each customer. Because this data is historical, it generally is not subject to revision or updates. High-performance storage savers reduce the cost of storing, managing, and processing this type of data. Organizations can archive tables or table partitions into the DB2 Analytics Accelerator to save space within the System z storage environment. All of that data is still maintained in the DB2 directory, and all of the queries that target that data are now directed only to the DB2 Analytics Accelerator. This approach dramatically reduces storage costs on System z and enables organizations to substantially increase the amount of history that is maintained for each subject area.

Connectivity options and workload balance

There is great flexibility for DB2 systems. Depending on the situation, users can connect multiple DB2 systems to a single DB2 Analytics Accelerator to share the capacity. A single DB2 system can also connect to multiple DB2 Analytics Accelerators to load sharing and redundancy for high reliability. Users can also connect multiple DB2 systems to multiple DB2 Analytics Accelerators and try different configurations.

When a single DB2 system is connected to multiple DB2 Analytics Accelerators, DB2 automatically balances the workload across multiple accelerators based on the accelerator's capacity and length of the request queue. DB2 checks the use of every eligible accelerator and routes the query to the most optimal one. Figure 4 shows the workload balance between two accelerators.

Workload balancing across multiple accelerators

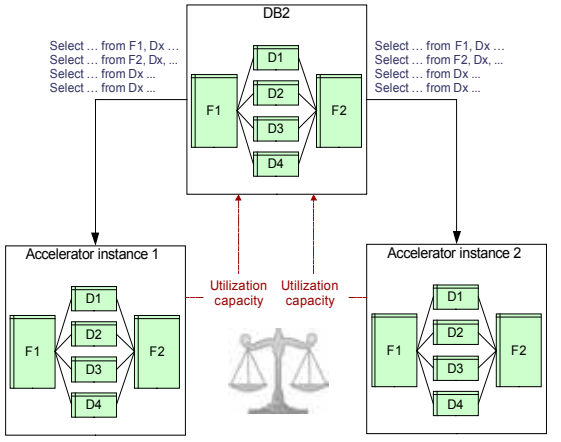

Figure 4: How DB2 distributes a workload across multiple accelerators

Disaster recovery

Disaster recovery can be extended to a DB2 Analytics Accelerator that is connected to DB2. This configuration has two data sharing group members, each of which is attached to an installation of DB2 Analytics Accelerator. Figure 5 shows how three applications are connected to one member and how they use DB2 Analytics Accelerator. Another two applications are on the other member.

Disaster recovery: Table loaded in one accelerator

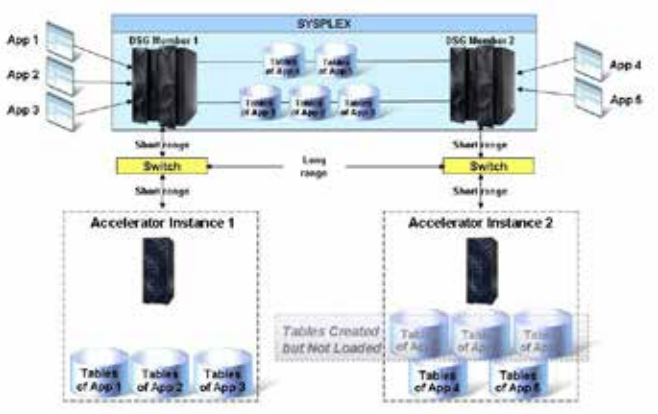

Figure 5: Application connections through DSG members

As shown in Figure 5, if the member 1 system is down, applications 1, 2 and 3 can reconnect to member 2. By redeploying tables for applications 1, 2, and 3 on accelerator instance 1, DB2 accelerator instance 2 can be used as the accelerator for them. Tables can also be loaded into multiple installations of DB2 Analytics Accelerator, and the system can be run with all of the applications. To shorten disaster recovery time, tables can be loaded redundantly into multiple installations of DB2 Analytics Accelerator.

DB2 integration

Deep integration of DB2 Analytics Accelerator with DB2 can affect many DB2 components. Some of the affected areas are more visible, and so more directly related to DB2 Analytics Accelerator:

- Optimizer and routing criteria
- Distributed data facility (DDF) or DRDA
- System parameters
- Special register
- EXPLAIN function
- Dynamic statement cache
- Instrumentation
- DB2 commands
- Administrative stored procedures

Query acceleration criteria

Query acceleration operates based on three criteria:

- *Environment and setup.* DB2 system, DRDA protocol, and query package
- *The query itself.* INSERT from select statement or read-only SELECT statement; references only tables that are deployed in the same accelerator. The SQL functionality that is required to execute the query is supported by the DB2 Analytics Accelerator. For static queries, the query acceleration is determined and fixed at package bind time. Tables must be defined to an accelerator and enabled for acceleration before the package is bound. Additionally, the accelerator must be active and started when a static query runs.
- *Apparent benefits*. DB2 determines whether acceleration speeds up the query based on the heuristics and the estimated query cost.

Query acceleration control knobs

Query acceleration is controlled by the special registers CURRENT QUERY ACCELERATION and CURRENT GET_ACCEL_ARCHIVE. Special register CURRENT QUERY ACCELERATION controls when DB2 sends queries to the DB2 accelerator and what DB2 does if the Accelerator server fails. It takes five possible values: NONE, ENABLE, ENABLE WITH FAILBACK, ELIGIBLE and ALL. To enable query acceleration, CURRENT QUERY ACCELERATION must be set as a value other than NONE. The special register CURRENT GET_ACCEL_ARCHIVE specifies whether a query that references an archived table uses the archived data. The possible values are YES and NO.

The special registers value can be set in three different scopes:

- System scope: zparms QUERY_ACCELERATION and GET_ACCEL_ARCHIVE set the default value of the special register at system level
- Statement scope: special registers CURRENT QUERY ACCELERATION and CURRENT GET_ACCEL_ARCHIVE can be set to control an individual SQL statement
- Application scope: the bind option can be specified for JDBC and ODBC applications, and BIND PACKAGE can be qualified by any identifier that is supported by DB2 profile tables

EXPLAIN function

The DB2 EXPLAIN function is enhanced to provide basic information about accelerator use. It indicates whether a query qualifies for acceleration and, if not, why it is not qualified. The access path details that are associated with the query are provided independently of DB2 EXPLAIN by DB2 Analytics Accelerator Studio.

When a query is accelerated, for each query, regardless of the number of query blocks that a query contains, the whole query has one row in PLAN_TABLE and DSN_QUERYINFO_TABLE. The PLAN_TABLE column ACCESSTYPE is A. The DSN_QUERYINFO_TABLE column QI_DATA contains the converted accelerator query text. If the query is not accelerated, REASON_CODE and QI_DATA columns provide details about why the query is not accelerated.

The EXPLAIN tables can be populated with information even if no accelerator is connected to DB2. Specifying EXPLAINONLY on the START ACCEL command does not establish any communications with an actual accelerator, but enables DB2 to consider its presence in the access path selection process. This virtual accelerator feature is useful when the resource limit facility (RLF) limits the execution of some long-running queries and users want to determine whether these long-running queries qualify for acceleration.

Accelerator Modeling

To help evaluate query workloads for acceleration on DB2 Analytics Accelerator and potentially save both elapsed time and processor usage, DB2 can provide accelerator modeling of query workloads. This modeling is performed with a new system parameter, ACCELMODEL, and new accounting fields, IFCID 3 and IFCID 148, which are similar to the ZIP and ZAP accounting fields.

Having DB2 Analytics Accelerator is not necessary to perform this accelerator modeling. Only queries deemed eligible for acceleration are included in these new accounting fields.

- QWAC_ACCEL_ELIG_ELA CL8
- The accumulated elapsed time that is spent processing SQL in DB2 that might be eligible for execution on an accelerator
- QWAC_ACCEL_ELIG_CP CL8
- The accumulated processing time that is spent on SQL in DB2 that might be eligible to run on an accelerator
- QWAC_ACCEL_ELIG_SE DS CL8

The accumulated processing time that is consumed on an IBM specialty engine while processing SQL in DB2 that might be eligible to run on an accelerator

For dynamic queries, accelerator modeling is enabled for a query during a full PREPARE, and elapsed time and processing time are accumulated while the query is run. During the full PREPARE, DB2 evaluates the dynamic query for acceleration. For dynamic queries already in the DB2 Dynamic Statement Cache (DSC), accelerator modeling is not enabled for those queries until the next full PREPARE that inserts the query into the DSC. Using RUNSTATS to invalidate the queries in the DSC before requesting accelerator modeling for workloads is recommended.

For static queries, accelerator modeling is enabled for a query during BIND/REBIND PACKAGE or auto-bind, and elapsed time and processing time are accumulated while the query runs. During the BIND/REBIND PACKAGE, DB2 evaluates the static query for acceleration. A new BIND/REBIND PACKAGE for packages that contain the static queries is needed after setting new zparm ACCELMODEL=YES. The existing bind option APREUSE(ERROR) can be specified to avoid access path changes. With APREUSE(ERROR), DB2 validates that the access path does not change during the REBIND. If the access path does change, then DB2 fails the REBIND.

Figure 6 shows the accelerator modeling account fields reported by IBM OMEGAMON® Performance Expert (PE).

Accelerator modeling as reported by OMPE

Figure 6: A breakdown of accounting fields as reported by OMEGAMON PE.

Accelerator-related instrumentation

To help performance monitoring, charge-back, capacity planning, and problem determination, instrumentation was enhanced to include accelerator-related information. The accelerator statistics, such as the total number of successful queries that are sent to the accelerator, total number of failed queries and query queue time, are collected in the Q8ST section of IFCID 2. The accelerator accounting data, such as query processing time and elapsed time that is spent in DB2 and in the accelerator, are collected in the Q8AC section of IFCID 3. Further illustration is provided in Figure 7, which is a sample of a STATISTICS REPORT from IBM OMEGAMON.

Sample statistics report for IBM DB2 Analytics Accelerator V4 using OMPE V5.2

Figure 7: A sample statistics report from OMEGAMON PE

DB2 Analytics Accelerator administrative stored procedures

DB2 Analytics Accelerator administrative stored procedures provide functions that are related to tables, queries, and accelerators. They can be invoked from IBM DB2 Analytics Accelerator Studio, run from the command line, or embedded in custom applications. The DB2 Analytics Accelerator stored procedures and related DB2 objects are created during the DB2 Analytics Accelerator installation.

Table 1 is an example list of DB2 Analytics Accelerator administrative stored procedures. Most of these stored procedures can be invoked from applications to automate some tasks, such as refreshing DB2 Analytics Accelerator data after ETL load.

Stored procedure	Description
ACCEL_ADD_ACCELERATOR	Pairs an accelerator with a DB2 subsystem
ACCEL_REMOVE_ACCELERATOR	Removes an accelerator from a DB2 subsystem and cleans up resources on the accelerator
ACCEL_ADD_TABLES	Adds a set of tables to the accelerator
ACCEL_ALTER_TABLES	Alters table definitions for a set of tables on the accelerator, only distribution and organizing keys
ACCEL_REMOVE_TABLES	Removes a set of tables from the accelerator
ACCEL_GET_TABLES_INFO	Lists a set of tables on the accelerator with detailed information
ACCEL_LOAD_TABLES	Loads data from DB2 into a set of tables on the accelerator
ACCEL_SET_TABLES_ACCELERATION	Enables or disables a set of tables for query offloading
ACCEL_GET_QUERY_EXPLAIN	Generates and retrieves the DB2 accelerator explain output for a query
ACCEL_GET_QUERIES	Retrieves active information, history query information or both from the accelerator
ACCEL_SET_TABLES_REPLICATION	Enables or disables incremental updates for one or more tables on an accelerator
ACCEL_GET_TABLES_DETAILS	Collects information about a set of tables regarding data changes and consistency, or moves operations with the High Performance Storage Saver
ACCEL_ARCHIVE_TABLES	Moves table partitions from DB2 for z/OS to a storage saver on an accelerator

Table 1: Example list of DB2 Analytics Accelerator stored procedures

DB2 Analytics Accelerator DB2 Command

Three DB2 commands were introduced to manage accelerator processing:

- DISPLAY ACCEL displays status information about accelerator servers
- START ACCEL activates the specified accelerator server that is stopped
- STOP ACCEL deactivates the specified accelerator server that is active and prevents the accelerator server from processing queries

Figure 8 is an example of the DISPLAY ACCELERATOR command. It shows the accelerator status.

```
DSNX810I  @ DSNX8CMD DISPLAY ACCEL FOLLOWS -
DSNX830I  @ DSNX8CDA 549
ACCELERATOR                      MEMB  STATUS  REQUESTS ACTV QUED MAXQ
-------------------------------- ----- -------- -------- ---- ---- ----
ZV3IDAA1                         VA1A STARTED       1662    0    0    0
LOCATION=ZV4IDAA2 HEALTHY
DETAIL STATISTICS
  LEVEL  = AQT03012
  STATUS = ONLINE
  FAILED QUERY REQUESTS                                =       0
  AVERAGE QUEUE WAIT                                   =    2084 MS
  MAXIMUM QUEUE WAIT                                   =   19312 MS
  TOTAL NUMBER OF PROCESSORS                           =       4
  AVERAGE CPU UTILIZATION ON COORDINATOR NODES         =    4.00%
  AVERAGE CPU UTILIZATION ON WORKER NODES              =     .00%
  NUMBER OF ACTIVE WORKER NODES                        =       1
  TOTAL DISK STORAGE AVAILABLE                         =   97713 MB
  TOTAL DISK STORAGE IN USE                            =       0 MB
DISPLAY ACCEL REPORT COMPLETE
DSN9022I  @ DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION
```

Figure 8: Example of the output from the DISPLAY ACCELERATOR command

Performance considerations

For query acceleration, it is important for users to keep aware of many factors:

- Consider trade-offs when determining which workload or queries to offload. Weigh the speed-up factor and processing savings against query volume for maximum throughput.
- Keep DB2 table and index statistics current so that DB2 can make optimal DB2 Analytics Accelerator offloading decisions.
- Watch for queries that return large result sets and push down data aggregation into the accelerator as applicable

When users load data to PureData System for Analytics, they can perform different tasks:

- Tune the AQT_MAX_UNLOAD_IN_PARALLEL WLM environment variable for the DB2 Analytics Accelerator load stored procedure. For optimal load performance, weigh the available system processor resources and the number of optimal concurrent active threads on PureData System for Analytics. A maximum of 10 threads is recommended.
- Before loading the tables into PureData System for Analytics from the DB2 Analytics Accelerator client, specify the appropriate distribution and organizing keys for tables

Powered by Netezza

DB2 Analytics Accelerator uses IBM Netezza technology as the accelerator. The PureData System for Analytics has a revolutionary design that is based on principles that allow it to provide an excellent price-to-performance ratio.

Four key components make up the PureData System for Analytics:

- SMP hosts
- Snippet blades, called S-Blades
- Disk enclosures
- Network fabric, which is not shown in Figure 9

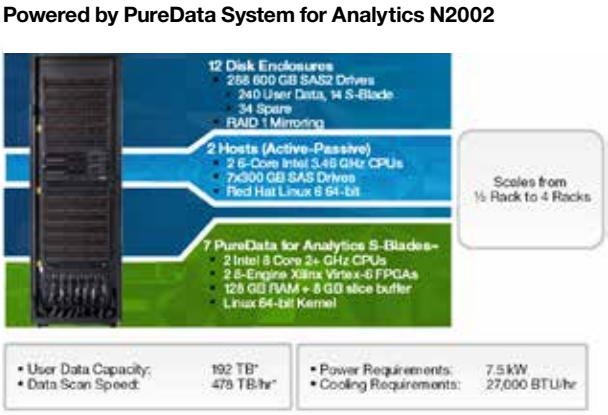

Figure 9: Key components of the PureData System for Analytics

The disk enclosures contain high-density, high-performance disks that are RAID protected. Each disk contains a slice of the data in the database table, along with a mirror of the data on another disk. The storage arrays are connected to the S-Blades through high-speed interconnections that allow all the disks to simultaneously stream data to the S-Blades at the fastest rate possible.

The DB2 Analytics Accelerator server also runs on the SMP host. SMP hosts are high-performance Linux servers that are set up in an active-passive configuration for high availability. The active host presents a standardized interface to external tools and applications such as BI and ETL tools and load utilities.

PureData System for Analytics compiles SQL queries into executable code segments, called snippets, creates optimized query plans, and distributes the snippets to the S-Blades for execution. S-Blades are intelligent processing nodes that make up the turbocharged message processing platform (MPP) engine of the appliance.

All system components are connected through a high-speed network fabric. PureData System for Analytics runs a customized IP-based protocol that fully uses the total cross-sectional bandwidth of the fabric and eliminates congestion even under sustained, intermittent network traffic.

The network is optimized to scale to more than a thousand nodes, while allowing each node to initiate large data transfers to every other node simultaneously. All system components are redundant.

While the hosts are active-passive, all other components in the appliance are hot-swappable. User data is fully mirrored, enabling better than 99.99 percent availability.

The PureData System for Analytics S-Blade, based on Netezza technology

Each S-Blade is an independent server that contains powerful multi-core processors, the Netezza unique multi-engine field programmable gate arrays (FPGAs), and gigabytes of RAM. All of these components are balanced and work concurrently to deliver peak performance.

FPGAs are commodity chips that are designed to process data streams at fast rates. Netezza employs these chips to filter out extraneous data based on the SELECT and WHERE clauses in the SQL statement, as quickly as data can be streamed off the disk. The process of data filtering reduces the amount of data by 95 – 98 percent, freeing up downstream components from processing unnecessary amounts of data.

The S-Blades also run an array of different database primitives such as sorts, joins, and aggregations in the processor cores. The processor cores are designed with ample room to run embedded algorithms of arbitrary complexity against large data streams for advanced analytics applications.

All key Netezza functions are performed on the S-Blade. Each S-Blade is a combination of a standard blade server and a database accelerator card that is provided by Netezza. It uses the IBM sidecar technology to easily combine the two blades to make them act as a single logical and physical entity. IBM commonly uses sidecar technology to expand its blade servers to add more memory or I/O blades to each server.

Applying data stream processing to DB2 queries

A key component of how Netezza performs is its streaming architecture processes data. The Netezza architecture uniquely uses the FPGA as a turbocharger, which is a huge performance accelerator. It allows the system to keep up with the data stream and accelerates the data stream through compression before processing it at line rates, ensuring no bottlenecks are in the I/O path.

Users can think of the way that data streaming works in Netezza as similar to an assembly line. The Netezza assembly line has various stages in the FPGA and processor cores. Each of these stages, along with the disk and network, operate concurrently, processing different chunks of the data stream at any point in time.

The concurrency within each data stream further increases performance as it relates to other architectures. Compressed data is streamed from disk onto the assembly line at the fastest rate that the physics of the disk allow. The data can also be cached, in which case it is served directly from memory instead of disk.

The first stage in the assembly, the Compress Engine within the FPGA core, picks up the data block and decompresses it at wire speed, transforming each block on disk into 4 – 8 blocks in memory. The result is a significant speedup of the slowest component in any data warehouse — the disk.

The disk block is then passed on to the Project engine or stage, which filters out columns based on parameters that are specified in the SELECT clause of the SQL query that is being processed. The assembly line then moves the data block to the Restrict engine, which strips off rows that are not necessary to process the query based on restrictions that are specified in the WHERE clause.

The Visibility engine also feeds parameters to the Restrict engine to filter out rows that should not be visible to a query — for example, rows that belong to a transaction that is not yet committed. The Visibility engine is critical in maintaining atomicity, consistency, isolation, and durability (ACID) compliance at streaming speeds in the Netezza unit.

The processor core picks up the decompressed, filtered data block and performs fundamental database operations, such as sorts, joins, and aggregations, on it. It also applies complex algorithms that are embedded in the snippet code for advanced analytics processing. It finally assembles all of the intermediate results from the entire data stream and produces a result for the snippet. The result is then sent over the network fabric to other S-Blades or the host, as directed by the snippet code.

Improved disk failover

Disk failover and resiliency are highly improved. Each disk is divided into three partitions: one that holds a slice of the user's data, a mirror of data on another disk, and a temp partition that is used to hold intermediate results. All of these partitions are mirrored, including the temp partition. The primary partition is mirrored in pairs in a RAID 1 format. The temp partition is laid out across a set of eight drives in the RAID 1+0 format, striped on mirrors.

Drive redistribution for S-Blade failures

All drives are visible to all S-Blades within a chassis. Therefore, if an S-Blade fails, its drives are redistributed among the remaining online S-Blades within a chassis. Each chassis has six S-Blades in it.

Comparison of PureData System for Analytics appliance models

Figure 10 compares the PureData System for Analytics appliance models. They range from a quarter-rack to 10 full-rack cabinets.

IBM DB2 Analytics Accelerator supports all models

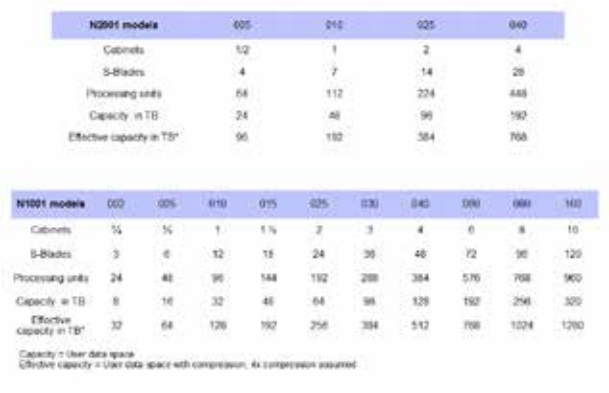

N2001 models	005	010	025	040
Cabinets	1/2	1	2	4
S-Blades	4	7	14	28
Processing units	56	112	224	448
Capacity in TB	24	48	96	192
Effective capacity in TB*	96	192	384	768

N1001 models	002	005	010	015	025	030	040	060	080	100
Cabinets	¼	½	1	1 ½	2	3	4	6	8	10
S-Blades	3	6	12	18	24	36	48	72	96	120
Processing units	24	48	96	144	192	288	384	576	768	960
Capacity in TB	8	16	32	48	64	96	128	192	256	320
Effective capacity in TB*	32	64	128	192	256	384	512	768	1024	1280

Capacity = User data space
Effective capacity = User data space with compression, 4x compression assumed

Figure 10: PureData System for Analytics appliance model comparison

Summary

IBM DB2 Analytics Accelerator brings fast performance to data-intensive and complex DB2 queries for data warehouse, business intelligence, and analytics workloads. DB2 Analytics Accelerator enables these queries to be transacted up to hundreds of times faster than was previously possible. Plus, the high performance and low cost of the DB2 Analytics Accelerator makes it ideal for organizations to use with data on the IBM System z platform.

For more information

For more information about IBM DB2 Analytics Accelerator, contact your IBM sales representative or go to **ibm.com/**software/data/db2/zos/analytics-accelerator.

About the authors

Namik Hrle works in the IBM Boeblingen Development Laboratory and is the lab's chief Information Management technologist, responsible for strategy and technology directions. As an IBM Distinguished Engineer and a member of the IBM Academy of Technology, Namik belongs to a small circle of the top technical leaders whose work and expertise affect the direction of IBM. He is a member of the Information Management Architecture Board, Software Group Architecture Board Steering Committee, Technical Experts Council, and many other IBM expert teams that work on strategic technology topics and address clients' information technology needs and requirements. He holds numerous patents, outstanding technical achievements, author recognition, and corporate awards.

Ruiping Li is a Senior Engineer in DB2 for z/OS development at IBM Silicon Valley Lab. She is the technical lead responsible for the DB2 query acceleration support for DB2 Analytics Accelerator. She was the development lead for DB2 V9 for z/OS optimistic locking support and DB2 V10 for z/OS time stamp with time zone data type support. She was also the key developer for other important functions, including IBM pureXML® index exploitation, complex queries, MQTs, and multiple CCSID features in DB2 for z/OS.

Wolfgang Hengstler has held positions within IBM for more than 30 years. He first worked in software development and then worked in product and market management for IBM Software Group, IBM Storage & Technology Group, IBM Tivoli®, and IBM Global Services. Wolfgang's projects have involved operating system components, OO wrapping technology, system automation products, and hosting services. He currently works in the IBM Information Management development lab in Boeblingen, Germany, and is part of a global product management team that focuses on data warehousing on IBM System z.

Acknowledgements

Thanks to the following people for their invaluable contributions:

- Surekha Parekh
- IBM Worldwide Marketing Program Director, DB2 for z/OS
- United Kingdom
- Wolfgang Hengstler
- Product Manager, IBM DB2 Analytics Accelerator
- Germany
- Guogen Zhang
- Former IBM Distinguished Engineer
- Silicon Valley Lab

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
April 2014

IMW14685-USEN-01