Revolution in IBM DB2 performance: IBM DB2 Analytics Accelerator

Namik Hrle, Ruiping Li and Wolfgang Hengstler, IBM

Contents

Short abstract

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 white paper answers the questions: what is business analytics, why is it important to your organization and how does the IBM solution differ from others? Then it will 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. The paper also describes which query types qualify for the new "access path" and 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, the paper 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 it needs 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 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, 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 them to use in their 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:

- Want to undertake a new reporting initiative on IBM System z to gain more insights
- Want to consolidate disparate data to their existing System z platform while benefiting from integrated operational BI
- Want 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 the 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 their 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 that 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:

- 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
- Multiple applications used for corporate reporting and business analysis
- Administration and management required for multiple platforms and complex data integration processes
- Have identified the value of consolidating data into a single easily managed platform (integrated OLTP and Analysis/ OLAP), but may have some concerns as to how analytic and traditional business intelligence workloads may perform on the mainframe
- The time it takes to deploy new data marts within the organization is 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 co-exists with many of their operational applications. The organization wishes to improve the performance of their 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 business intelligence 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:

- Query performance and execution time of individual queries or overall workloads could be improved significantly, freeing up millions of instructions per second (MIPS) and storage space, therefore reducing processing cost
- Ability to execute queries which were either forgotten or blocked previously by the administrator due to performance issues
- Increased organization agility by being able to more rapidly respond with immediate, accurate information and deliver new insights to business users
- Reporting is consolidated on System z to 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

This chart describes an IBM DB2 system, including applications, tools, and DB2 itself. Inside DB2 we have some familiar components, such as Data Manager, Buffer Manager, Log Manager, Internal Release Lock Manager, 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.

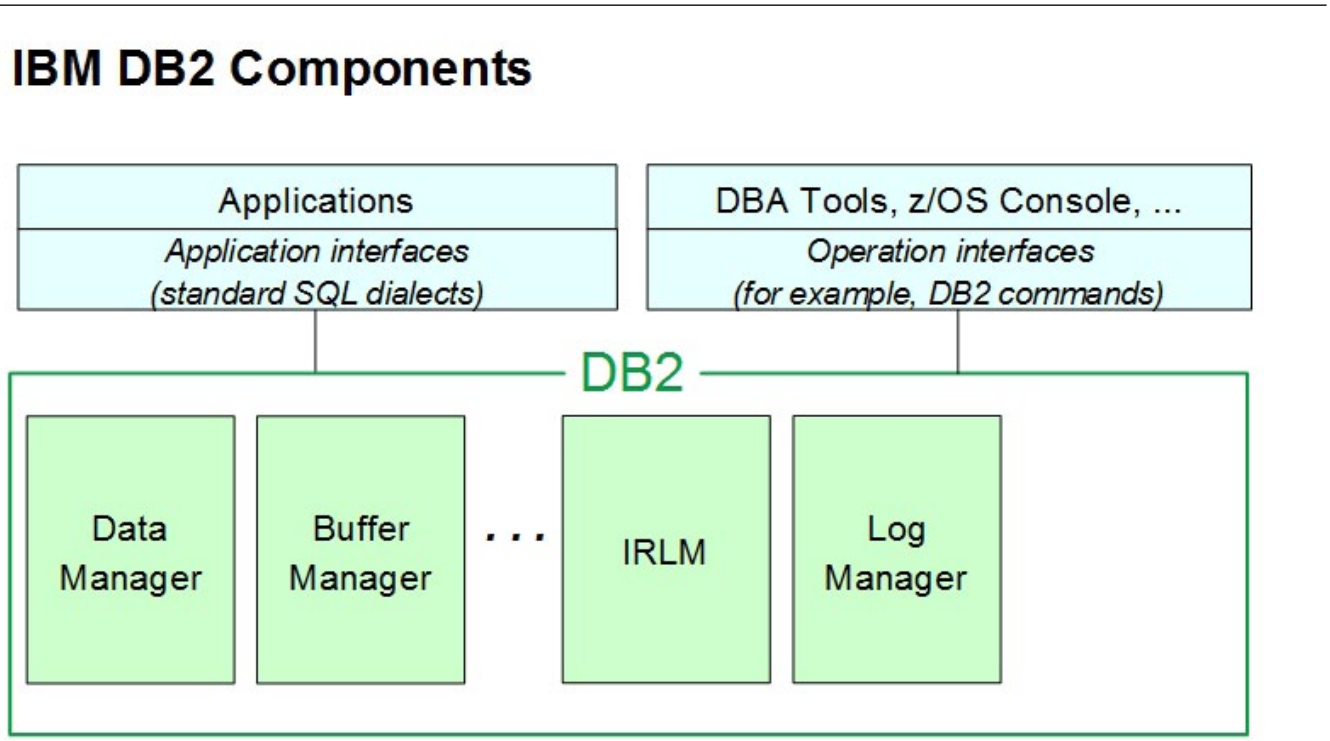

The following chart illustrates that IBM now has a new "virtual" DB2 component, called IBM DB2 Analytics Accelerator. DB2 Analytics Accelerator has hardware and software components, based on IBM Netezza technology, used to accelerate complex queries that typically are seen in analytics applications.

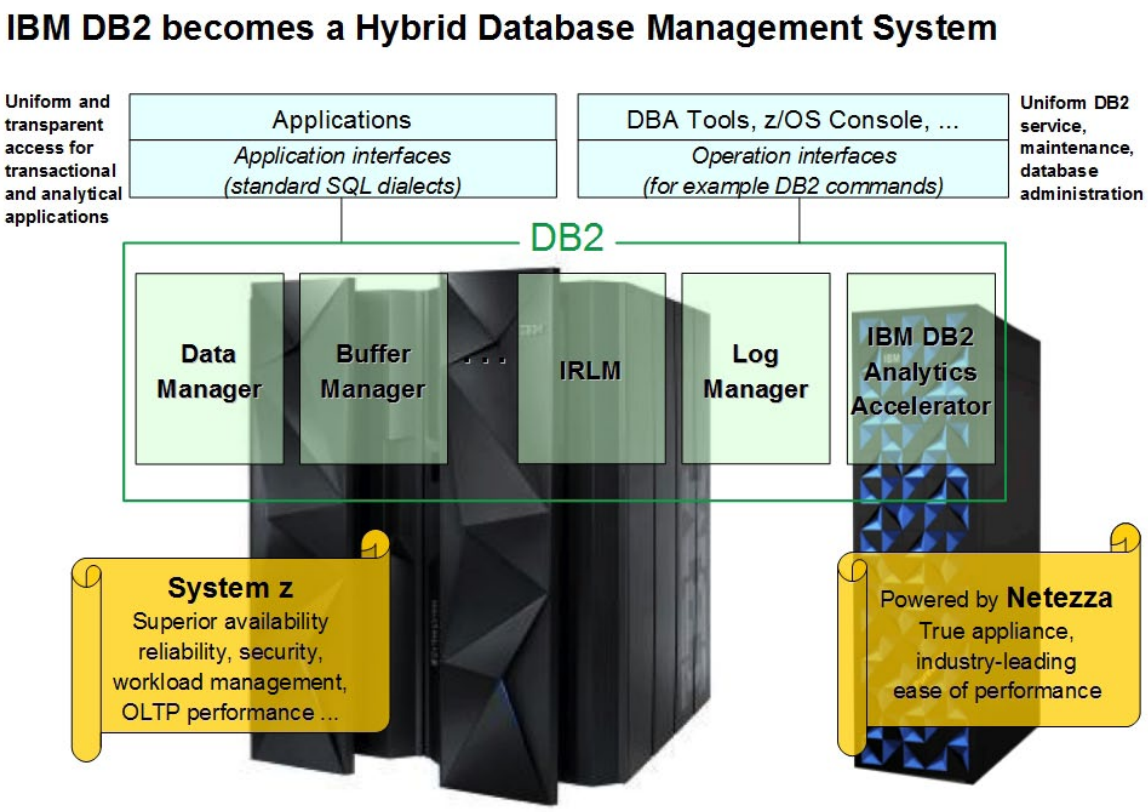

The IBM PureData™ System for Analytics appliance, connected to DB2, is enhanced to act as a DB2 accelerator. And DB2 has been enhanced with query acceleration to execute queries in DB2 Analytics Accelerator. It provides lightning-fast query performance transparently to the DB2 applications—at an affordable 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. PureData System for Analytics appliance provides superb data warehouse performance and ease of use of an appliance.

Using it as an accelerator, it's not necessary to tend to administrative processes as with a stand-alone 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. This is true, as you can see in the EXPLAIN output.

Query execution process flow

The following chart 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 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 will convert the query into Netezza syntax, and route 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 PureData System for Analytics Appliance box and sends the result back to DB2 through DRDA. And the result is sent out to the application.

Also shown on the following chart are the heartbeat messages from DB2 Analytics Accelerator to DB2, with DB2 Analytics Accelerator availability and performance indicators.

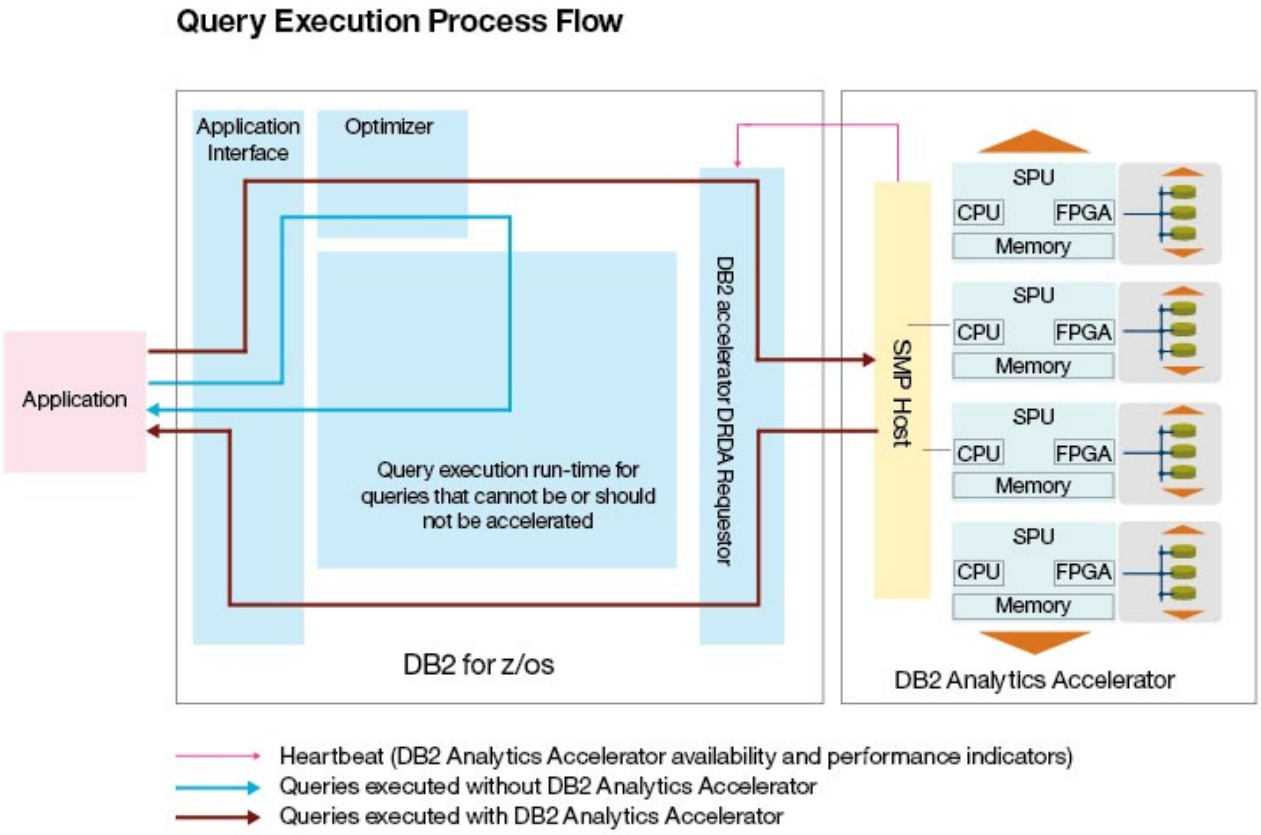

DB2 Analytics Accelerator content maintenance

Query acceleration by DB2 Analytics Accelerator is a new access path for DB2, just like an index access path. In order 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 load the data from DB2 into DB2 Analytics Accelerator afterwards, and you can refresh the data content periodically. It uses the DB2 UNLOAD utility to unload data, in parallel, to USS pipes. DB2 Analytics Accelerator will read the data and convert it into LOAD to PureData System for Analytics Appliance.

Partitions belonging to the same table can be loaded in parallel, with user-controlled degree of parallelism to balance CPU and load throughput. Refresh or updates are done on a "per table," "per partition" or an 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 are 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, enabling decisions to 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 integrated appliance form factor of the Accelerator.

DB2 Analytics Accelerator table definition and deployment

Before loading data, you need to define tables to be accelerated. All administrative tasks are achieved through DB2 stored procedures for DB2 Analytics Accelerator. DB2 Analytics Accelerator Studio provides GUI interface to these stored procedures and convenience in performing the administrative tasks. Applications can invoke the stored procedures directly.

The stored procedures will update the pseudo catalog tables, which provide the necessary information to support the query acceleration. Defining and deploying tables into DB2 Analytics Accelerator is simple. You simply identify the tables for which queries need to be accelerated, and load data and enable it for query acceleration.

High-performance storage saver

Most analytical systems are based on data in which over 95 percent of it is 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 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 select those tables or table partitions and not require them to consume space within the System z storage environment. All of the data is still maintained in the DB2 directory, and all the queries that target that data are now only directed to the Accelerator. This not only dramatically reduces storage costs on System z, it also enables organization to substantially increase the amount of history maintained for each subject area.

Connectivity options

Depending on your situation, you can connect multiple DB2 systems to a single DB2 Analytics Accelerator to share the capacity. Or a single DB2 system can connect to multiple DB2 Analytics Accelerators to load sharing and redundancy for high reliability. You can also connect multiple DB2 systems to multiple DB2 Analytics Accelerators and "mix and match" them.

There is great flexibility for DB2 systems.

Disaster recovery considerations

Disaster recovery can be extended to DB2 Analytics Accelerator connected to DB2. This is a configuration with two data-sharing group members, each attached with a DB2 Analytics Accelerator. The following diagram shows how applications 1, 2 and 3 are connected to member 1, and are using DB2 Analytics Accelerator 1. Applications 4 and 5 are on the other member.

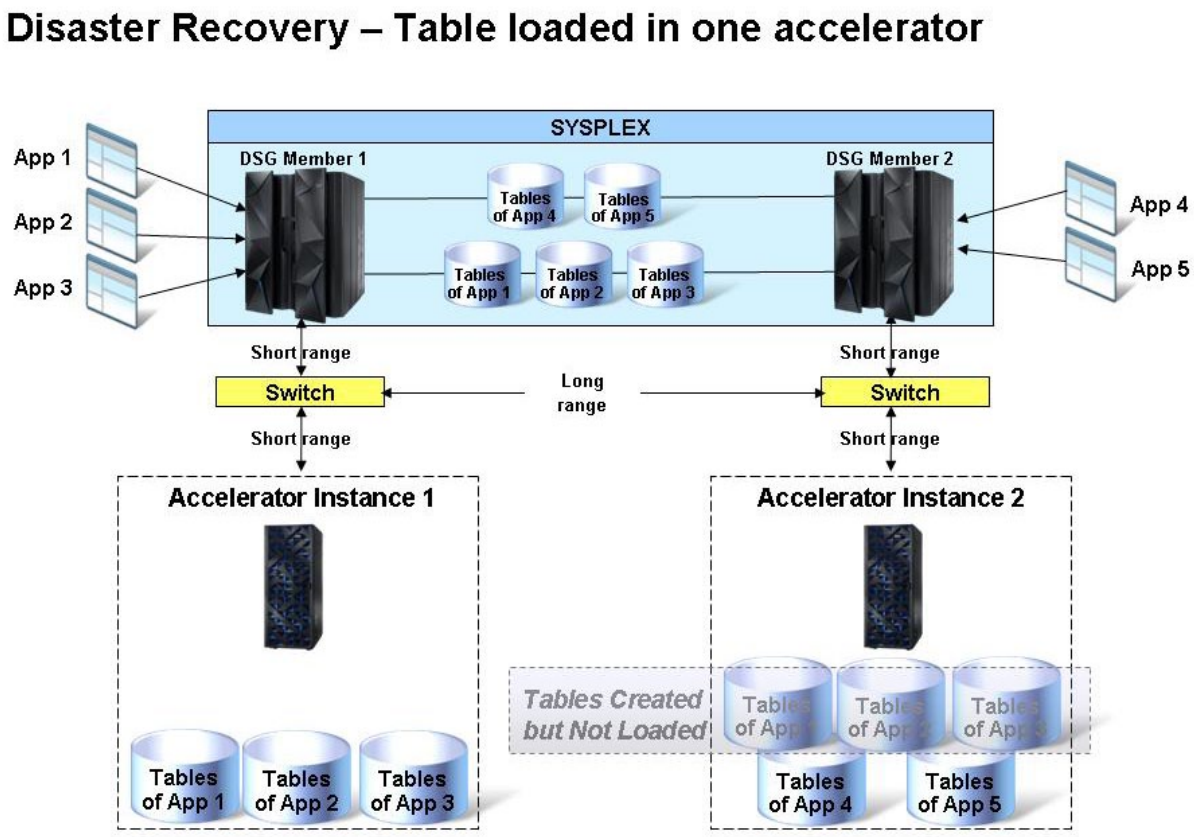

In case member 1 system is down, applications 1, 2 and 3 can reconnect to member 2. By redeploying (loading and enabling) tables for applications 1, 2, and 3 on DB2 Analytics Accelerator 2, DB2 Analytics Accelerator 2 can be used as the accelerator for them. Tables can be loaded into multiple DB2 Analytics Accelerators. And the system can be up and running with all the applications. To shorten disaster recovery time, tables can be loaded redundantly into multiple DB2 Analytics Accelerators. Then only enabling is needed for the switch.

DB2 integration

Deep integration of DB2 Analytics Accelerator with DB2 has affected many DB2 components, and this list identifies some of the more visible areas related to DB2 Analytics Accelerator:

- Optimizer and routing criteria
- Distributed data facility (DDF)/DRDA
- System parameters
- Special register
- Explain
- Dynamic statement cache
- Instrumentation
- DB2 commands
- DB2 Analytics Accelerator administrative stored procedures

Query acceleration criteria

There are three parts in the criteria:

1. Environment/setup: DB2 system, DRDA protocol, query package.
2. The query itself: INSERT from select statement or dynamic read-only SELECT statement; references only tables that are deployed in the same accelerator; the SQL functionality required to execute the query is supported by the DB2 Analytics Accelerator.
3. DB2 checks that acceleration will speed up the query based on the heuristics and the estimated query cost.

Query acceleration control knobs

Query acceleration is controlled by both the zparm parameters and a new special register, CURRENT QUERY ACCELERATION. Following is a list of three zparm parameters and the special register setting for the query acceleration. The zparm parameters are for the system level. The new special register can be used to control an individual SQL statement. It inherits the default value from the zparm QUERY_ACCELERATION, and can override the zparm setting.

- ACCEL zparm—specifies whether accelerator servers can be used with the DB2 subsystem, and how the accelerator servers are to be enabled and started. An accelerator server cannot be started unless it is enabled.

Value	Description
NO	The accelerator servers cannot be used with the DB2 subsystem.
AUTO	The accelerator servers are automatically enabled and started when the DB2 subsystem is enabled.
COMMAND	The accelerator servers are automatically enabled when the DB2 subsystem is started. The accelerator servers can be started with the DB2 START ACCEL command.

To enable query acceleration, ACCEL should be set as AUTO or COMMAND.

- ACCEL_LEVEL zparm—specifies which version of the accelerator that DB2 is to use (only required for DB2 Version 9.1).

Value	Description
V1	The accelerator servers are to use IBM Smart Analytics Optimizer Version 1. V1 is the default.
V2	The accelerator servers are to use DB2 Analytics Accelerator.

To enable query acceleration, ACCEL_LEVEL should be set as V2.

- QUERY_ACCELERATION zparm—determines the default value that is to be used for the CURRENT QUERY ACCELERATION special register.

Value	Description
NONE	No query is accelerated. NONE is the default.
ENABLE	The queries are accelerated only if DB2 determines that it is advantageous to do so. If there is an accelerator failure while a query is running, or the accelerator returns an error, DB2 returns a negative SQLCODE to the application.
ENABLE_ WITH_FAIL-BACK	The queries are accelerated only if DB2 determines that it is advantageous to do so. If the accelerator returns an error during the PREPARE of first OPEN for the query, DB2 executes the query without the accelerator. If the accelerator returns an error during a FETCH or a subsequent OPEN, DB2 returns the error to the user, and does not execute the query.
ELIGIBLE	The queries are accelerated if they are eligible for acceleration. DB2 does not use cost information to determine whether to accelerate the queries. Queries that are not eligible for acceleration are executed by DB2. If there is an accelerator failure while a query is running, or the accelerator returns an error, DB2 returns a negative SQLCODE to the application.
ALL	The queries are accelerated if they are eligible for acceleration. DB2 does not use cost information to determine whether to accelerate the queries. Queries that are not eligible for acceleration are not executed by DB2, and an SQL error is returned. If there is an accelerator failure while a query is running, or the accelerator returns an error, DB2 returns a negative SQLCODE to the application.

- CURRENT QUERY ACCELERATION special register—the special register is used to override the QUERY_ACCELERATION zparm. CURRENT QUERY ACCELERATION special register determines whether SQL queries are considered for query acceleration, and what DB2 does if the accelerator server fails. There are five possible values: NONE, ENABLE, ENABLE_WITH_ FAILBACK, ELIGIBLE and ALL (as described earlier).

To enable query acceleration, CURRENT QUERY ACCELERATION should be set as a value other than NONE.

EXPLAIN function

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

When a query is accelerated, for each query, regardless the number of query blocks that a query contains, the whole query has one row in both PLAN_TABLE and DSN_ QUERYINFO_TABLE. PLAN_TABLE column ACCESSTYPE is 'A'. 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 the details why the query is not accelerated.

Note that the EXPLAIN tables can be populated with above described information even if there is no accelerator connected to DB2. Specifying EXPLAINONLY on 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 is the virtual accelerator feature. It is useful when resource limit facility (RLF) limits the execution of some long-running queries and you want to find out whether these long-running queries qualify for acceleration.

This lists the columns in DSN_QUERYINFO_TABLE, and the column names shown in red font contain information more specific to DB2 Analytics Accelerator enhancements.

Column Name	Column Contents
QUERYNO	A number that identifies the statement that is being explained.
QBLOCKNO	A number that identifies each query block within a query.
QINAME1	If REASON_CODE = 0, the name of the accelerator
QINAME2	If REASON_CODE = 0, the location of the accelerator
APPLNAME	The name of the application plan for the row. Applies only to embedded EXPLAIN statements that are executed from a plan or to statements that are explained when binding a plan. A blank indicates that the column is not applicable.
PROGNAME	The name of the program or package containing the statement being explained. Applies only to embedded EXPLAIN statements and to statements explained as the result of binding a plan or package. A blank indicates that the column is not applicable.
VERSION	The version identifier for the package. Applies only to an embedded EXPLAIN statement executed from a package or to a statement that is explained when binding a package. A blank indicates that the column is not applicable.
COLLID	The collection ID for the package. Applies only to an embedded EXPLAIN statement that is executed from a package or to a statement that is explained when binding a package. A blank indicates that the column is not applicable.
GROUP_MEMBER	The member name of the DB2 that executed EXPLAIN. The column is blank for non-data sharing.
SECTNOI	The section number of the statement.
SEQNO	The sequence number in the table.
EXPLAIN_TIME	The time at which the statement is processed. This time is the same as the BIND_TIME column in PLAN_TABLE.
TYPE	The type of the output for this row. “A” means that this row is for a query that DB2 attempts to run on an accelerator server. The value in column REASON_CODE indicates the outcome.
REASON_CODE	If 0, the query qualifies for acceleration. Otherwise, the query cannot be accelerated. More details on the next chart.
QI_DATA	If REASON_CODE = 0, the text of the converted SQL statement (sent to DB2 Analytics Accelerator). Otherwise, the description of the reason for not qualifying for acceleration.
SERVICE_INFO	IBM internal use only
QB_INFO_ROWID	IBM internal use only

This is a list of the REASON_CODE values to tell why a query is not qualified for acceleration.

Value	Description
0	Query qualifies for acceleration
1	No active accelerator was found when EXPLAIN was executed.
2	The special register CURRENT QUERY ACCELERATION is set to NONE.
3	The query is a DB2 short running query or re-routing to the accelerator is not considered advantageous.
4	The query is not read-only
5	The query is running under the private protocol.
6	The cursor is defined as scrollable or rowset cursor.
7	The query refers to multiple encoding schemes.
8	The query FROM clause specifies a data-change-table-reference.
9	The query contains a correlated table expression.
10	The query contains a recursive common table expression reference.
11	The query contains an unsupported expression. QI_DATA contains the expression text.
12	The query references table table-name that is either not defined in accelerator, or the table is defined, but is not enabled for query re-routing.
13	The accelerator accelerator-name containing the tables of the query is not started.
14	The column column-name referenced in the query is altered in DB2 after the data is loaded in the accelerator.
15	The query references a DB2 10 new SQL feature.
16	The query is not from a package.
900 through 999	IBM internal use

System-scope instrumentation

Instrumentation has been enhanced to include accelerator-related information. Here you see a sample of the STATISTICS REPORT from IBM OMEGAMON®.

```
1   LOCATION: PMOV91A                OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V5R1)                PAGE: 1-23
      GROUP: N/P                              STATISTICS REPORT - LONG                REQUESTED FROM: NOT SPECIFIED
     MEMBER: N/P                                                                               TO: NOT SPECIFIED
  SUBSYSTEM: V91A                                                                   INTERVAL FROM: 09/06/11 21:49:41.35
DB2 VERSION: V9                                  SCOPE: MEMBER                                  TO: 09/06/11 23:41:50.70

---- HIGHLIGHTS ------------------------------------------------------------------------------------------------
INTERVAL START  : 09/06/11 21:49:41.35   SAMPLING START: 09/06/11 21:49:41.35   TOTAL THREADS      :   76.00
INTERVAL END    : 09/06/11 23:41:50.70   SAMPLING END  : 09/06/11 23:41:50.70   TOTAL COMMITS      :  109.00
INTERVAL ELAPSED:       1:50:52.248034   OUTAGE ELAPSED:          1:17.097273   DATA SHARING MEMBER:    N/A

ZGRYPHON          ACCELERATOR            QUANTITY   ZGRYPHON          CONTINUED            QUANTITY
---------------------------  ------------------   ---------------------------  ------------------
CONNECTS                                    2.00   AVG QRY QUEUE LEN (3 HRS)                 0.00
REQUESTS                                    9.00   AVG QRY QUEUE LEN (24 HRS)                0.00
REQUESTS TIMED OUT                          0.00   HWM QRY QUEUE LENGTH                      0.00
REQUESTS FAILED                             0.00   DATA SKEW                                 0.00
BYTES SENT                               4630.00   AVG QUEUE WAIT ELAPSED TIME            0.000000
BYTES RECEIVED                         224887.00   MAX QUEUE WAIT ELAPSED TIME            0.000000
MESSAGES SENT                              27.00
MESSAGES RECEIVED                          27.00   PROCESSING CAPACITY                       0.00
BLOCKS SENT                                 0.00   PROCESSORS                                1.62
BLOCKS RECEIVED                             5.00
ROWS SENT                                   0.00   QUERY REQUESTS SUCCESSFUL                 1.00
ROWS RECEIVED                               0.00   QUERY REQUESTS FAILED                     1.00
                                                     QUERY REQUESTS INVALID                    0.00
SVCS TCP/IP ELAPSED TIME                7.036035
ACCELERATOR CPU TIME                    0.000000   SHR MEM WORKER NODES   (MB)               0.00
ACCELERATOR ELAPSED TIME                0.000001     AVG IN USE           (MB)               0.00
ACCELERATOR WAIT TIME                   0.000000     MAX IN USE           (MB)               0.00

CUR ACTIVE REQUESTS                         0.00   DISK STORAGE AVAILABLE (MB)            98842.63
MAX ACTIVE REQUESTS                         0.00     IN USE               (MB)               0.81
                                                     IN USE FOR DB        (MB)               0.81
                                                   DATA SLICES                              3.25

                                                   MEM COORD AVG IN USE   (MB)               0.00
                                                   MEM WORKER AVG IN USE  (MB)               0.00
```

Thread-scope instrumentation

This shows the accounting report for user ADMF001 plan db2jcc_a and the accelerator XGRYPHON activity.

```
   LOCATION: PMOV91A                  OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V5R1)                PAGE: 1-7
      GROUP: N/P                            ACCOUNTING REPORT - LONG                  REQUESTED FROM: NOT SPECIFIED
     MEMBER: N/P                                                                                  TO: NOT SPECIFIED
  SUBSYSTEM: V91A                            ORDER: PRIMAUTH-PLANNAME                 INTERVAL FROM: 09/06/11 21:52:08.31
DB2 VERSION: V9                                 SCOPE: MEMBER                                     TO: 09/06/11 23:43:26.65

PRIMAUTH: ADMF001  PLANNAME: db2jcc_a

---- DISTRIBUTED ACTIVITY ------------------------------------------------------------------------------------------------
SERVER              : ZGRYPHON            CONVERSATIONS INITIATED:     1.00    #COMMT(1)SENT:       0    MESSAGES SENT     :      14.33
PRODUCT ID          : AQT                 #CONVERSATIONS QUEUED  :        0    #ROLLB(1)SENT:       0    MESSAGES RECEIVED:      14.33
METHOD              : DRDA PROTOCOL       CONVERSATION TERMINATED:      N/A    SQL SENT      :    6.33    BYTES SENT        :    2403.33
REQUESTER ELAP.TIME:           2.402285   #RLUP THREADS          :      N/A    ROWS RECEIVED:    0.00    BYTES RECEIVED    : 138728.33
SERVER ELAPSED TIME:                N/A                                                                 BLOCKS RECEIVED   :       3.33
SERVER CPU TIME     :                N/A
DBAT WAITING TIME   :                N/A
#DDF ACCESSES       :                  3

#COMMIT(2) SENT     :                  0   #BACKOUT(2) SENT       :        0    #BKOUT(2) R.R:       0    #LASTAGN.SENT     :          0
                                           SUCCESSFULLY ALLOC.CONV:      N/A    TRANSACT.SENT:    1.00    STMT BOUND AT SER:        N/A
                                           MAX OPEN CONVERSATIONS :      N/A    MSG.IN BUFFER:    0.00    #FORGET RECEIVED  :          0
                                           #CONT->LIM.BL.FTCH SWCH:      N/A    #PREPARE SENT:       0
                                           #COMMIT(2) RESP.RECV.  :        0

ACCELERATOR  IDENTIFIER                      ACCELERATOR          TOTAL VALUE    TOTAL TIME      AVERAGE VALUE  AVERAGE TIME
-----------  ------------------------------  -------------------  -----------  ----------------  -------------
PRODUCT      AQT02010                        OCCURRENCES                    3                             1.00
SERVER       ZGRYPHON                        CONNECTS                       3                             1.00
                                             REQUESTS                      16                             5.33
                                              TIMED OUT                     0                             0.00
                                              FAILED                        0                             0.00
                                             SENT
                                              BYTES                      7210                          2403.33
                                              MESSAGES                     43                            14.33
                                              BLOCKS                        0                             0.00
                                              ROWS                          0                             0.00
                                             RECEIVED
                                              BYTES                    416185                        138728.33
                                              MESSAGES                     43                            14.33
                                              BLOCKS                       10                             3.33
                                              ROWS                          0                             0.00

                                             ELAPSED TIME
                                              SVCS TCP/IP                          7.206857                        2.402286
                                              ACCUM ACCEL                          0.000061                        0.000020
                                             CPU TIME
                                              SVCS TCP/IP                          0.015023                        0.005008
                                              ACCUM ACCEL                          0.000000                        0.000000
                                             WAIT TIME
                                              ACCUM ACCEL                          0.000000                        0.000000
```

DB2 Analytics Accelerator administrative stored procedures

Here you see a list of DB2 Analytics Accelerator administrative stored procedures. Most of these stored procedures can be invoked from applications to automate some tasks, such as refresh DB2 Analytics Accelerator data after ETL load.

IBM DB2 Analytics Accelerator Administrative Stored Procedures	
ACCEL_ADD_ACCELERATOR	Pairing an accelerator to a DB2 subsystem
ACCEL_TEST_CONNECTION	Check of the connectivity from DB2 procedures to the accelerator
ACCEL_REMOVE_ACCELERATOR	Removing an accelerator from a DB2 subsystem and cleanup resources on accelerator
ACCEL_UPDATE_CREDENTIALS	Renewing the credentials (authentication token) in the accelerator
ACCEL_ADD_TABLES	Add a set of tables to the accelerator
ACCEL_ALTER_TABLES	Alter table definitions for a set of tables on the accelerator (only distribution and organizing keys)
ACCEL_REMOVE_TABLES	Remove a set of tables from the accelerator
ACCEL_GET_TABLES_INFO	List set of tables on the accelerator together with detail information
ACCEL_LOAD_TABLES	Load data from DB2 into a set of tables on the accelerator
ACCEL_SET_TABLES_ACCELERATION	Enable or disable a set of tables for query off-loading
ACCEL_CONTROL_ACCELERATOR	Controlling the accelerator tracing, collecting trace and detail of the accelerator (software level etc.)
ACCEL_UPDATE_SOFTWARE	Update software on the accelerator (transfer versioned software packages or apply an already transferred package, new: also list software both on z/OS and accelerator side)
ACCEL_GET_QUERY_DETAILS	Retrieve statement text and the query plan for a running or completed query that is routed to the DB2 accelerator
ACCEL_GET_QUERY_EXPLAIN	Generate and retrieve the DB2 accelerator explain output for a query
ACCEL_GET_QUERIES	Retrieve active and/or history query information from 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 with regard to data changes (consistency) or move 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

DISPLAY ACCELerator command

This is an example of the -DISPLAY ACCELERATOR command. It shows a group scope status.

DISPLAY ACCELerator

```
-DIS ACC(*) SCOPE(GROUP)
DSNX830I ) DSNX8CMD DISPLAY ACCELERATOR FOLLOWS -
ACCELERATOR                      MEMB   STATUS REQUESTS ACTV QUED MAXQ
-------------------------------- ---- -------- -------- ---- ---- ----
BLINK1                           DB1A  STARTED    32769    2    5   23
BLINK1                           DB1B  STARTED    23456    1    0    2
BLINK1                           DB1C  STARTED      734    0    0    4
BLINK1                           DB1D  STARTED     9210    0    0    1
BLINK2                           DB1A  STOPPED    37235    1    7   17
BLINK2                           DB1B  STOPPED       47    0    0    0
BLINK2                           DB1C  STARTED        2    0    0    0
BLINK2                           DB1D  STOPPED        0    0    0    0
BLINK3                           DB1A  STARTED     3256    5   23   41
BLINK3                           DB1B  STOPPED       92    0    0    2
BLINK3                           DB1C  STARTED       87    0    0    7
BLINK3                           DB1D  STOPPED       21    0   11   11
DISPLAY ACCELERATOR REPORT COMPLETE
DSN9022I ) DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION
```

This shows an example of the display of the active accelerator BLINK1 for DB1D.

DISPLAY ACCELerator

```
-DIS ACC(BLINK1) LIST(ACTIVE) SCOPE(LOCAL) MEMBER(DB1D)

DSNX810I ) DSNX8CMD DISPLAY ACCEL FOLLOWS -
DSNX830I ) DSNX8CDA
ACCELERATOR                      MEMB   STATUS REQUESTS ACTV QUED MAXQ
-------------------------------- ---- -------- -------- ---- ---- ----
BLINK1                           DB1D  STARTED     9210    7    5    9

LOCATION=ACCELERATOR1 HEALTHY

DETAIL STATISTICS
LEVEL = AQT02010
STATUS = ONLINE
FAILED QUERY REQUESTS = 3
AVERAGE QUEUE WAIT = 99
MAXIMUM QUEUE WAIT = 400
TOTAL NUMBER OF PROCESSORS = 4
AVERAGE CPU UTILIZATION ON COORDINATOR NODES = 45.00%
AVERAGE CPU UTILIZATION ON WORKER NODES = 40.00%
NUMBER OF ACTIVE WORKER NODES = 2
TOTAL DISK STORAGE AVAILABLE = 93000 MB
TOTAL DISK STORAGE IN USE = 56100 MB
DISK STORAGE IN USE FOR DATABASE = 36100 MB
DISPLAY ACCEL REPORT COMPLETE
DSN9022I ) DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION
```

Performance considerations

Query acceleration:

- Consider trade-offs when determining which workload or queries to offload. Speed up factor and CPU savings will need to be weighed against query volume for maximum throughput.
- Keep DB2 table and index statistics up to date so that DB2 could make optimal DB2 Analytics Accelerator offloading decisions
- Watch for queries that return large result sets and push down data aggregation into accelerator as applicable

Loading data to PureData System for Analytics:

- Tune the AQT_MAX_UNLOAD_IN_PARALLEL WLM environment variable for the DB2 Analytics Accelerator load stored procedure and weigh the available the system CPU resources and number of optimal concurrent active threads (recommended maximum of ten threads) on PureData System for Analytics for optimal load performance.
- Specify appropriate distribution and organizing keys for tables before loading the tables into PureData System for Analytics from DB2 Analytics Accelerator client.

Powered by Netezza

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

There are four key components that make up the PureData System for Analytics (see illustration at bottom of page):

- SMP hosts
- Snippet blades (called S-Blades)
- Disk enclosures
- Network fabric (not shown in the illustration)

Accelerator powered by IBM PureData System for Analytics N1001

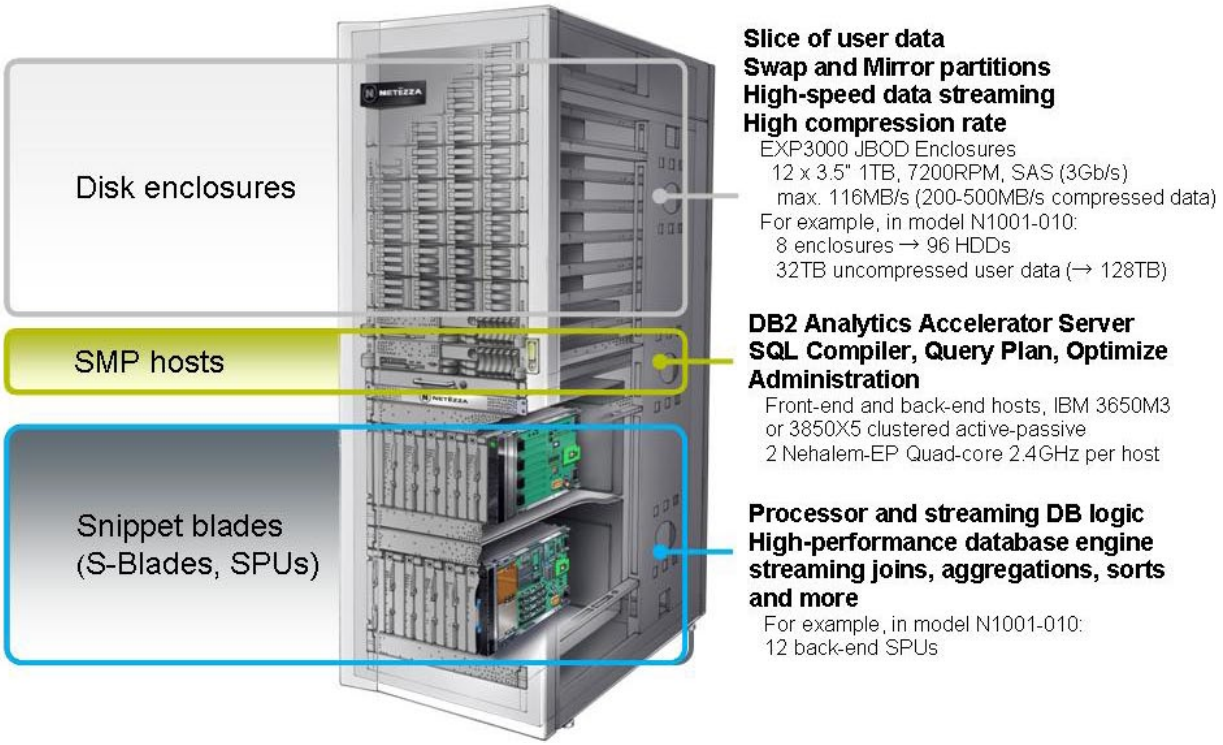

The disk enclosures contain high-density, high-performance disks that are RAID protected. Each disk contains a slice of the data in the database table, along with a mirror of the data on another disk. The storage arrays are connected to the S-Blades through high-speed interconnects that allow all the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

The 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 business intelligence (BI) and extract, transform and load (ETL) tools and load utilities.

It compiles SQL queries into executable code segments called snippets, creates optimized query plans and distributes the snippets to the S-Blades for execution. The DB2 Analytics Accelerator server also runs on the SMP host. S-Blades are intelligent processing nodes that make up the turbocharged message processing platform (MPP) engine of the appliance.

Network Fabric: All system components are connected by means of 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 CPUs, Netezza's unique multi-engine field programmable gate arrays (FPGAs) and gigabytes of RAM — all balanced and working concurrently to deliver peak performance.

FPGAs are commodity chips that are designed to process data streams at extremely 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 to 98 percent, freeing up downstream components from processing unnecessary amounts of data.

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

The S-Blade is where the key Netezza functions are performed. Each S-Blade is a combination of a standard blade server and a database accelerator card provided by Netezza. It uses IBM's "sidecar" technology to easily combine the two blades to make them act as a single logical and physical entity.

NOTE: The sidecar technology is commonly used by IBM to expand their blade servers to add more memory or I/O blades to each server.

Applying data stream processing to DB2 queries

A key component of Netezza's performance is the way in which its streaming architecture processes data. The Netezza architecture uniquely uses the FPGA as a turbocharger—a huge performance accelerator that not only allows the system to keep up with the data stream, but to actually accelerate the data stream through compression before processing it at line rates, ensuring no bottlenecks in the I/O path.

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

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

The first stage in the assembly line, the Compress Engine within the FPGA core, picks up the data block and decompresses it at wire speed, instantly transforming each block on disk into from four to eight 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 specified in the SELECT clause of the SQL query 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 specified in the WHERE clause.

The Visibility engine also feeds in additional parameters to the Restrict engine, to filter out rows that should not be "seen" by a query—for example, rows belonging to a transaction that is not committed yet. 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 uncompressed, 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 the intermediate results together 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.

IBM PureData for Analytics N1001 Appliance models

Powered by Netezza

Model	002	005	010	015	025	030	040	060	080	100
Cabinets	1/4	1/2	1	1 1/2	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 (TB)	8	16	32	48	64	96	128	192	256	320
Effective Capacity (TB)*	32	64	128	192	256	384	512	768	1024	1280

IBM DB2 Analytics Accelerator supports all the models

Capacity = User data space
Effective Capacity = User data space with compression (4x compression assumed)

Comparing PureData System for Analytics appliance models

This chart compares the PureData System for Analytics Appliance models. They range from a quarter rack to ten full-rack cabinets.

Minimizing disk failures

Disk failover and resiliency is 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 RAID 1+0 format (striped on mirrors).

Minimizing S-Blade failures

All drives are visible to all S-Blades within a chassis. Thus, 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).

Summary

In this white paper, we've examined how IBM DB2 Analytics Accelerator brings extremely fast performance to data-intensive and complex DB2 queries for data warehouse, business intelligence and analytics workloads. It enables these queries to be transacted up to hundreds of times faster than was previously possible. The paper covered how IBM DB2 Analytics Accelerator makes this possible—and how the high performance and low cost of the IBM DB2 Analytics Accelerator makes it ideal for organizations to use it with data on the IBM System z platform.

In reading this paper, you have learned about the IBM DB2 Analytics Accelerator architecture and how to use the accelerator, control acceleration and maintain its content. The paper should give you a knowledge of the new access path and which workloads and query types apply for the new access type.

For more information

For more information on the 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, he 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 as well as addressing customers information technology needs and requirements. He is the holder of numerous patents, outstanding technical achievements and 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 the IBM DB2 Analytics Accelerator. She has been the development lead for DB2 9 for z/OS new feature optimistic locking support and DB2 10 for z/OS new feature timestamp with time zone data type support, and has also been the key developer for some other important functionality. This functionality includes pureXML index exploitation, complex queries, MQTs, and multiple CCSID features in DB2 for z/OS.

For more than 30 years, **Wolfgang Hengstler** has held positions within IBM—first in software development and then 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 focusing on data warehousing on IBM System z.

Acknowledgements

Thanks to the following people for their invaluable contributions to this paper and for making it possible.

- 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
January 2013