

Reliability and Performance with IBM DB2 Analytics Accelerator Version 4.1

IBM Redbooks Solution Guide

The IBM® DB2® Analytics Accelerator for IBM z/OS® (simply called *DB2 Accelerator* or just *Accelerator* for this IBM Redbooks® Solution Guide) is a high-performance appliance that integrates the IBM zEnterprise infrastructure and the IBM Netezza technology to accelerate data-intensive and complex queries in a DB2 for z/OS highly secure and available environment (see Figure 1). DB2 and the Accelerator appliance form a self-managing hybrid environment, concurrently and efficiently running online transaction processing and online transactional analytical processing together with business intelligence and online analytic processing workloads.



Figure 1. DB2 Analytics Accelerator

IBM DB2 Analytics Accelerator V4.1 expands the value of high-performance analytics. DB2 Analytics Accelerator V4.1 opens to static SQL applications and rowset processing, minimizes data movement, reduces latency, and improves availability.

The IBM Redbooks publication *Reliability and Performance with IBM DB2 Analytics Accelerator Version V4.1*, SG24-8213 helps technical decision makers understand the benefits of the Version 4.1 Accelerator when used with DB2 11 for z/OS. The installation of the new functions and the advantages to existing analytical processes as measured in a controlled test environment are described. The IBM DB2 Analytics Accelerator Loader for z/OS V1.1, a tool that facilitates the data population of the DB2 Accelerator, is also introduced.

Did you know?

By investing in the DB2 Accelerator, query elapsed times can be significantly reduced, CPU cost for query processing can be moved to the Accelerator, and storage cost for data that is no longer modified can be reduced by moving that data onto less costly disks attached to the Accelerator, all while protecting the IBM System z® investment.

The Accelerator has also been enhanced in recent releases to handle not only dynamic SQL but also static SQL. This improvement can extend query processing to the transactional and batch workloads for those queries that qualify for acceleration. The possibilities that this improvement brings to System z is almost endless.

Business value

“Wicked fast!” was the phrase used to describe the DB2 Accelerator by a customer during the implementation of a proof of concept (PoC).

The DB2 Accelerator, using Netezza technology, and combining with DB2 for z/OS creates a solution capable of increasing query processing by 2000X. It allows for the completion of analytics queries in seconds that used to take hours (see Figure 2). Most important, it allows queries that would never have completed because of their complexity to finally produce an answer, in only seconds or minutes.

MODE	Expressed in	CLASS 2 Elapsed	CLASS 2 CPU	CLASS 2 Elapsed Savings	CLASS 2 CPU Savings	% CLASS 2 Elapsed Savings	% CLASS 2 CPU Savings
DB2	seconds	35629	15719	33772	15698.5	94.79%	99.87%
Accelerator	seconds	1857	20.5				
DB2	minutes	593	262	562	261.66		
Accelerator	minutes	31	0.34				
DB2	hours	9.89	4.36	9.37	4.3543		
Accelerator	hours	0.52	0.0057				

Overall initial results across 25 queries, numbers include queries that performed better in DB2

Figure 2. DB2 to Accelerator query comparisons

In the Accelerator’s more recent release, it can process static SQL in addition to dynamic SQL. This allows it to be used for long running production SQL that in the past would have been relegated to batch and second and third shift. With the ability to complete an SQL in seconds or minutes rather than hours, it frees up resources and restriction of when SQL can be executed.

Although not its primary purpose, it can also reduce mainframe MIPS (millions of instructions per second) usage. If a query that ran for 2 hours for example on the mainframe is now executed on the Accelerator, not only does it complete in a significantly shorter elapsed time, the MIPS that would have been consumed by running that query on the mainframe are now available.

Solution overview

Business Analytics are no longer considered a “nice to have” with low performance expectations. Data warehousing and Business Analytics are becoming mission critical applications with the same quality of service (QoS) requirements as traditional Online Transaction Processing (OLTP) workloads. The applications must be reliable, available, and secure: attributes synonymous with System z. Sometimes, there is also a need to have the analytics workload coexist with traditional OLTP in a mixed workload environment. The ability to run mixed workloads is a major differentiator when determining a warehouse-hosting platform.

Today’s analytic solutions must incorporate the same ideals of all mainframe applications:

- Leverage your investment in IBM System z for data warehousing
- Transparently accelerate DB2 complex queries
- Implement highly available analytics

To best achieve these characteristics today, an analytics solution should be based on the DB2 Analytics Accelerator. The backbone of the DB2 Analytics Accelerator for z/OS is its use of the IBM PureData System for Analytics appliance, an appliance that is based on Netezza technology that delivers easy-to-use warehouse agility and performance. The IBM PureData System for Analytics appliance also enables the Accelerator to increase its overall data capacity to 32 TB of uncompressed raw data, with the ability to easily expand to a petabyte. However, when an IBM PureData System for Analytics appliance is implemented as part of DB2’s query accelerator, the IBM PureData System for Analytics appliance “loses” its identity by becoming a component available only for DB2 for z/OS use. This happens because the Accelerator appears as a virtual component of DB2 for z/OS.

The DB2 Analytic Accelerator is also part of the zEnterprise Analytics System 9700/9710 solution. The IBM zEnterprise Analytics System 9700/9710 are a set of fully integrated and scalable data warehouse and analytics solutions that are designed to give an organization the insight it needs to work smarter. They combine software, server, and storage resources to put the correct answers in the hands of the corporate decision makers today; placing your business in the best possible position to answer the questions of tomorrow.

The included components have been selected, tested, optimized, and priced to enable customers to deploy a next generation decision system. They have been sized for a wide variety of configurations, ranging from the modest of database sizes and users up through the petabyte systems with hundreds to tens of thousands of users.

Deeply optimized and ready to use, the zEnterprise Analytics System 9700/9710 can help quickly turn information into insight and deliver that insight where and when it is needed.

Solution architecture

DB2 consists of multiple resource managers (for example, Data Manager, Buffer Manager, Relational Data System, and so on) that perform its various functions. When the Accelerator is installed, it essentially becomes one more resource. The difference is that these other resource managers exist as part of the DB2 subsystem and run in DB2 while the Accelerator runs on the IBM PureData System for Analytics appliance. Figure 3 is an example of DB2 for z/OS with no Accelerator attached; Figure 4 is an example of DB2 with the addition of the DB2 Analytics Accelerator for z/OS.

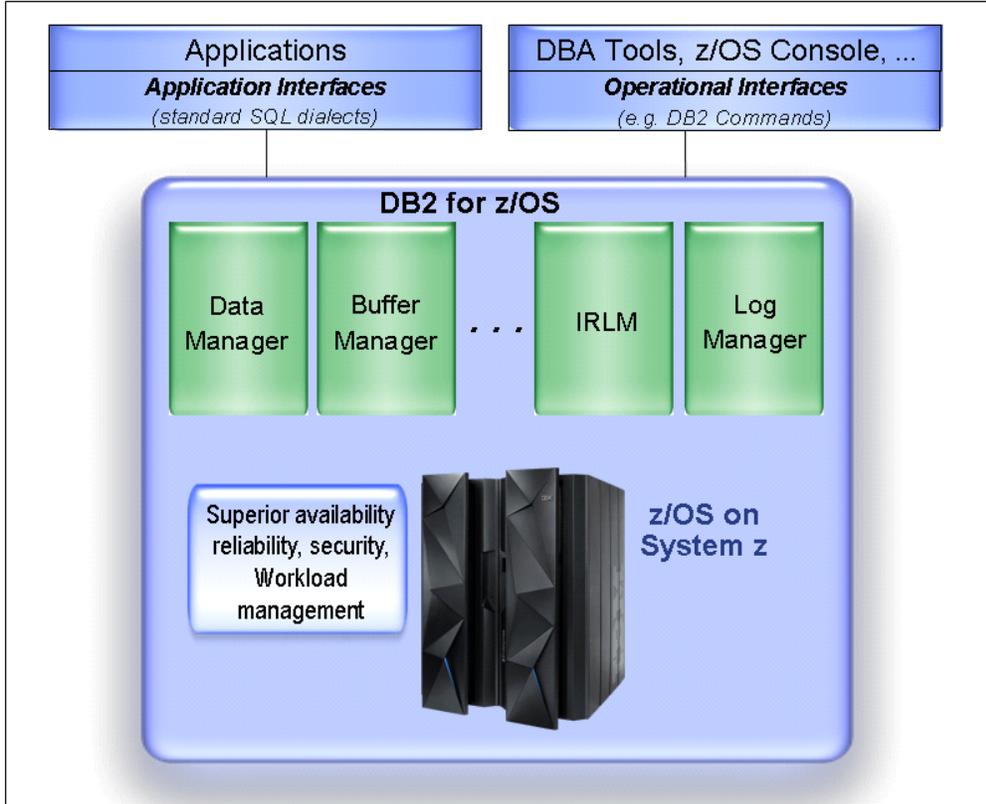


Figure 3. DB2 with no Accelerator attached

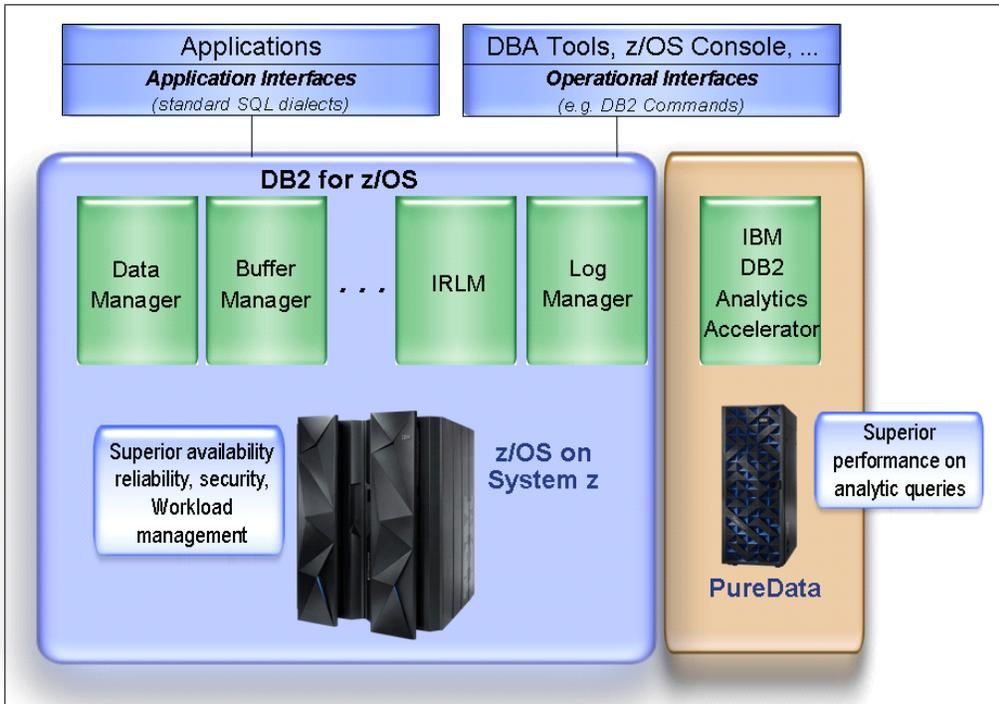


Figure 4. DB2 subsystem with attached DB2 Analytics Accelerator for z/OS

The Accelerator's subcomponent status is key to how the Accelerator processes a query. When a query arrives at DB2, the optimizer determines the best access path. With the Accelerator attached, the Accelerator becomes an extra access path. If the DB2 optimizer decides the query qualifies for offload to the Accelerator, the query is moved to the Accelerator and executed. Use of the Accelerator is transparent to any task, creating a dynamic or static SQL statement that might eventually qualify to run on the Accelerator. If DB2 decides the query will run in a more effective manner natively on DB2, then the query runs the way dynamic and static SQL has always run in DB2. Because the optimizer makes the routing decision, failure of the Accelerator is not a show stopper. The Accelerator is just no longer considered an access path resource and all queries execute natively in DB2 on z/OS without any outage. The decision of how to handle a failure at the Accelerator is controlled by a subsystem parameter (sets the default for the DB2 special register) or DB2 special register. This is the basis for no program changes being required to take advantage of the Accelerator; where to run the query is always the optimizer's decision (not the application's).

Some DB2 SQL query components and characteristics are incompatible with the Accelerator and prevent certain query types from ever being eligible to run on the Accelerator. Only read-only dynamic and static SQL statements qualify to run on the Accelerator. In addition, the SELECT statement can use a cursor that is defined as scrollable or for rowset processing (multi-row SELECT). Only the DB2 optimizer decides what queries are routed to the Accelerator.

The query data

Before DB2 can consider running a query on the Accelerator, a copy of the data necessary to satisfy the query must be preloaded to the Accelerator. Supplied DB2 stored procedures are used to define tables in the Accelerator and initiate data unload from DB2 and data load to the Accelerator. The unload and load processes can be started using either the DB2 Analytics Accelerator Studio, an Eclipse based GUI used for various administrative tasks that are hosted by the IBM DB2 Analytics Accelerator Studio V4.1 (this is the most current version as of this writing) client, or by explicitly invoking the DB2 stored procedures. Data is unloaded from a table or table partition to the UNIX System Services pipes and passed to the Accelerator for processing. Data in the Accelerator is historical and always "stale" to some extent. The Accelerator always contains a copy of the data that is stored in DB2. In most scenarios, data is never "moved" from DB2 to the Accelerator. The exception is when taking advantage of the Accelerator's high-performance storage save (HPSS). Also, the Accelerator provides a way to continually keep the Accelerator and DB2 in sync through incremental update by using IBM InfoSphere Change Data Capture for z/OS.

Using the supplied DB2 stored procedures, the Accelerator's data load process can load an entire table or just a partition of the table. It can add partitions to a table that already exists in the Accelerator, loading those added partitions, and deleting a partition's contents on the Accelerator. Loading data at the partition level is supported only for range-partitioned universal table spaces, which might be another reason for moving to universal table spaces.

Tables that are loaded by a single stored procedure are always loaded sequentially. However, parallelism is available for loading data into the Accelerator. If different stored procedure invocations start the load processes, those load tasks can run in parallel. Loading an entire partition table, or refreshing a subset of a table's partitions, those partitions can be processed in parallel. Parallel table load is supported only for universal table spaces. In some cases, the Accelerator can run SQL concurrently to a parallel data refresh.

Using the Accelerator

What needs to be accomplished for a query to have the option of running at the Accelerator? Besides a few APARs necessary to enable Accelerator processing, and the necessary software and hardware installations, the Accelerator must be enabled, data must be loaded, and the query must qualify (both the query type and the SQL functionality the query uses).

After the installation process is completed, you must set a few DB2 subsystem parameters that let the Accelerator run. One of those parameters sets the default value for the Accelerator special register, which can be set to completely disallow all use of the Accelerator, to route all eligible SQL to the Accelerator, or to route eligible SQL to the Accelerator with the option to run certain queries natively on DB2 if they fail to run on the Accelerator.

The Accelerator currently handles most commonly used data warehousing SQL features and will be continually improved.

The following major functions are delivered by the DB2 Accelerator:

- Dynamic and static query processing
- Load and incremental update
- Accelerator monitoring
- Accelerator maintenance procedures
- Running multiple workloads on one Accelerator
- High availability and balancing workloads
- Online data archiving
- DB2 Analytics Accelerator Loader for z/OS (charge based add-on product)

Usage scenarios

Looking at DB2 Analytics Accelerator usage scenarios is looking at where the transactional source data is being analyzed today.

- Rapid acceleration of business critical queries

Moving the business analytics to System z allows the corporate data to be stored, managed, and analyzed all within the single zEnterprise platform. With System z, the data and the analytics take advantage of a secure, reliable, scalable, and highly energy efficient platform design that can process multiple diverse workloads simultaneously. The overall customer solution results in performance improvements and cost reductions.

- Derive business insight from z/OS transaction systems

If the data is offloaded to a distributed data warehouse or data mart, the Accelerator might help to simplify and consolidate complex infrastructures and reduce time-consuming ETL processes. This results in lowering data latency, enhancing reliability, increasing security and decreasing overall total cost of ownership (TCO). The Accelerator is used to reduce IT sprawl for analytics.

- Reduce IT sprawl for analytics initiatives

If the data is not being analyzed yet, because queries are set aside as a result of performance or cost challenges or because a new BI applications is being considered, then the Accelerator can be attached to the mainframe to complement DB2 for z/OS, to form a hybrid infrastructure, that best hosts this new workload. The Accelerator is used to derive business insight from z/OS transaction systems.

- Improve access to historical data and lower storage costs

The DB2 Accelerator can deliver reduced storage cost by storing non-updated data on the less expensive DB2 Accelerator disk. This moves the data not being modified to less costly disk storage attached to the DB2 Accelerator but still leaves that data readily available for use by queries running on the DB2 Accelerator.

Integration

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

Supported platforms

The IBM DB2 Analytics Accelerator for z/OS runs on DB2 for z/OS (currently DB2 10 and DB2 11) environments and complies with their prerequisites in terms of System z resources.

Ordering information

You can request delivery for DB2 Analytics Accelerator for z/OS, V4.1 (product number 5697-DAB) through CBPDO and ServerPac. These customized offerings are for Internet delivery in countries where ShopzSeries product ordering is available. Internet delivery reduces software delivery time and allows you to install software without the need to handle tapes. For details about Internet delivery, see the Shopz help information:

<http://www.software.ibm.com/ShopzSeries>

You may also request a quote for IBM zEnterprise Analytics System 9700 and 9710 from the following location:

<http://bit.ly/1gZApTv>

Or, you may call (in the US) 1-877-426-3774; specify **Info Mgmt** as the priority code.

Related information

To find technical details about the DB2 Accelerator, see the following resources:

- *Reliability and Performance with IBM DB2 Analytics Accelerator Version 4.1*, SG24-8213
- *Hybrid Analytics Solution using IBM DB2 Analytics Accelerator for z/OS V3.1*, SG24-8151
- *Optimizing DB2 Queries with IBM DB2 Analytics Accelerator for z/OS*, SG24-8005
- White paper: *IBM DB2 Analytics Accelerator: A revolution in performance*, IMW14685USEN:
<http://bit.ly/U7li0v>
- IBM DB2 Analytics Accelerator for z/OS page:
<http://www-03.ibm.com/software/products/en/db2analacceforzos>
- IBM Analytics on zEnterprise page:
<http://www-01.ibm.com/software/os/systemz/badw/>
- IBM Offering Information page (to search for announcement letters, sales manuals, or both):
http://www.ibm.com/common/ssi/index.wss?request_locale=en

On this page, enter `DB2 Analytics Accelerator for z/OS`, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2014. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document was created or updated on June 10, 2014.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbook@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips1186.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml> The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Cognos®
DB2®
IBM®
IBM PureData™
InfoSphere®
PureData™
QMF™
Query Management Facility™
Redbooks®
Redbooks (logo)®
SPSS®
System z®
TM1®
z/OS®
zEnterprise®

The following terms are trademarks of other companies:

Netezza, and N logo are trademarks or registered trademarks of IBM International Group B.V., an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries. Other company, product, or service names may be trademarks or service marks of others.