

IBM DB2 with BLU Acceleration

IBM Redbooks Solution Guide

IBM DB2® with BLU Acceleration is a revolutionary technology, delivered in DB2 for Linux, UNIX, and Windows Release 10.5. Using *dynamic in-memory* columnar technologies, BLU Acceleration enables organizations to perform analytic queries at breakthrough speeds.

With query response times up to 100 times faster than earlier systems, BLU Acceleration provides a simple, fast, and easy-to-use solution for organizations that seek rapid access to business answers. Armed with this new information, organizations can lower costs and quickly gain an edge on competitors that are slower to draw insight from their data.

Yes, the industry is abuzz with discussion of in-memory columnar data processing. Yet BLU Acceleration offers much more. For example, it delivers significant improvements in database compression but does not require you to have all your data in memory.

Furthermore, BLU Acceleration is built directly into the DB2 kernel (see Figure 1). It is not just an additional feature; it is an integral part of DB2 and every part of DB2 is aware of it. BLU Acceleration still uses the same storage unit of pages, the same buffer pools, and the same backup and recovery mechanisms, so BLU Acceleration tables can coexist with traditional row tables in the same schema, storage, and memory, and you can query a mix of traditional row and BLU Acceleration tables at the same time.

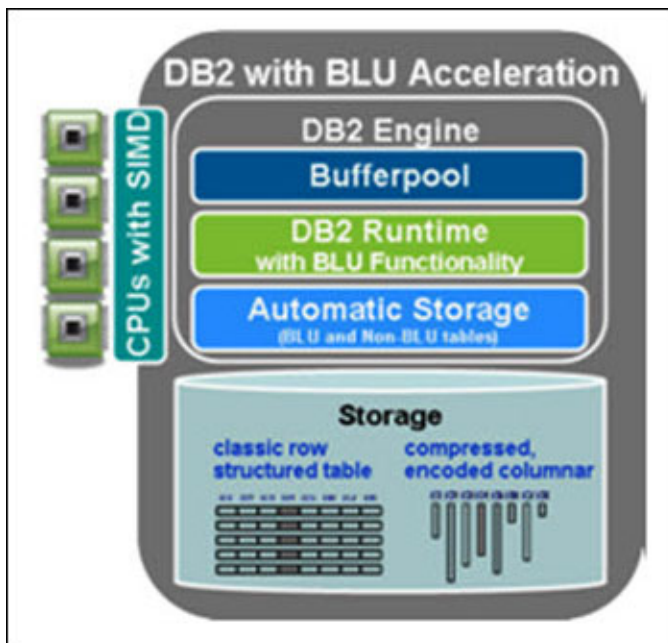


Figure 1. DB2 10.5 engine with BLU Acceleration

Did you know?

BLU Acceleration from IBM differs from solutions that other vendors offer in that it allows the unified computing of online transaction processing (OLTP) and analytics data inside a single database. This helps remove barriers and speed results for users.

Business value

Although the cost of memory continues to fall, data growth is exponential, so analytics systems that require all data to be kept in memory are no longer practical. This is where BLU Acceleration technology enters the picture, delivering in-memory performance even when the need for active data is bigger than the available memory space.

Extends existing investments

DB2 with BLU Acceleration runs on existing infrastructure, helping to extend the return on investments that the organization has already made. With relatively loose hardware requirements, system configurations can be optimized to meet specific cost targets and service levels. As data volumes increase, the cost of manipulating each piece of the data decreases dramatically.

Customers can choose between IBM POWER or x86 processor architectures. In addition, Advanced Workload Management capabilities enhance transactional and analytics operations while supporting more concurrent users at higher service levels than alternative solutions.

Proven business support

BLU Acceleration runs on DB2, so it builds upon the business-proven performance, reliability, and security of the platform while adding its own unique memory management and query optimization features.

DB2 is used around the world for online transactional processing (OLTP) and online analytic processing (OLAP) workloads, with decades of high-availability performance and a reputation for flexible scalability, both up and out. DB2 has consistently recorded top results in industry-standard benchmark tests, with multiple 100 TB-and-larger systems in production use every day.

Easy to adopt and use

The real breakthrough in BLU Acceleration is in how easily organizations can adopt the technology. Competing solutions often require migration to new platforms, database software, and business applications, making the transition both costly and resource intensive. But, because BLU Acceleration is integrated into DB2, you can move into it at your own pace.

The move to column-based tables can be accomplished one at a time or all at once. For guidance, DB2 includes tools that help IT personnel assess the impact of using BLU tables in different workloads and analytical environments.

Solution overview

In introducing BLU Acceleration, IBM did not just bolt a new component onto DB2. Support for the new solution already existed within the DB2 engine, so there is no new skill gap for DBAs and developers to overcome.

Simply stated, you need only to install DB2 10.5 on a supported platform and enable a single configuration parameter to tell DB2 that you want to optimize the default settings for analytical workload characteristics. After you create your database, create your tables, and load the needed data using common DB2 tools; your workloads immediately run faster and without the need for SQL syntax or schema changes.

At the center of BLU Acceleration is a column-organized table store combined with *actionable compression* that operates on a column and page level to save storage space. The column organization eliminates the need for creating and maintaining secondary indexes and aggregates. In DB2 10.5, both column-organized and traditional row-organized tables can coexist in the same database.

DB2 with BLU Acceleration includes several features that work together to make it a significant advancement in technology:

- **Simplicity and ease-of-use**

Purchasing new hardware for a BLU Acceleration deployment might not be necessary. Autonomics and intelligence is built into the DB2 engine to optimize your hardware for analytic workloads. The goal is to minimize deployment complexity and database maintenance tasks for DBAs.

Manual tuning efforts are minimized, replaced with a single registry variable (`DB2_WORKLOAD=ANALYTICS`) that automates initial tuning of an analytics database environment. With this registry variable set, the database manager and database parameters are automatically tuned by the DB2 engine to optimize analytic-type workloads.

BLU Acceleration is simple to implement and easy to use. BLU Acceleration tables coexist with traditional row-organized tables in the same database, the same table spaces, and the same buffer pools. All you have to do is create your BLU Accelerated tables, load the data into tables, and run your queries.

- **Column store**

BLU Acceleration uses dynamic in-memory columnar technology. Each column is physically stored in a separate set of data pages.

This column-organized approach brings many benefits to analytic workloads, which compared to other workloads tend to involve more table joins, grouping, aggregates, and so on, and often only access a subset of columns in a query. By storing data in column-organized tables, DB2 only needs to pull the wanted column data into memory rather than the entire row. As a result, more column data values can be packed into processor cache and memory, significantly reducing I/O from disks.

Column-organized tables typically compress much more effectively. BLU Acceleration compresses data by column; the probability of finding repeating patterns on a page increases significantly when the data on the page is from the same column. In addition, being able to store both row-organized and column-organized tables in the same database allows you to use BLU Acceleration even in database environments that require mixed OLTP and OLAP workloads.

- Adaptive compression

A key aspect of DB2 with BLU Acceleration is the manner in which data is encoded on disk, which enables significant compression and, because DB2 uses a different encoding scheme, allows data to be scanned or compared *while it remains encoded*. The ability to scan compressed data allows DB2 to delay materialization until absolutely necessary, and do as much work as possible without decoding the data. This, in turn, conserves processing power and I/O throughput.

BLU Acceleration stores data using column organization, so each column is compressed with its own compression dictionaries. When BLU Acceleration compresses data, it uses traditional adaptive mechanisms plus a form of Huffman encoding, in which data is compressed based on the frequency of values in a column (so values that appear many times are compressed more than other values that do not appear as often). For instance, if a column displays the states of the United States, the relative populations of the states provides the likelihood that California will be listed many more times than Vermont. So DB2 can encode California with a short, single bit (1), and encode Vermont with eight bits (11011001). For even greater compression, DB2 combines this feature with prefix encoding and offset coding for better compression in various scenarios.

In addition, because BLU Acceleration handles query predicates without decoding the values, more data can be packed in processor cache and buffer pools. The result is less disk I/O, better use of memory, and more effective processor utilization. Query performance is better and storage needs are significantly reduced.

- Parallel vector processing

Single-instruction, multiple data (SIMD) is a common processor technology. When SIMD-enabled hardware is used, DB2 with BLU Acceleration has special algorithms to take advantage of the built-in parallelism of the processors. In this way, you can use special hardware instructions that work on multiple data elements with a single instruction.

For instance, without SIMD, for a query that requires a predicate comparison of sales orders sent in December 2013 (see Figure 2), the data elements must be processed one at a time, with multiple iterations needed to process a single instruction. With BLU Acceleration using both SIMD processor parallelism and efficient columnar compression, multiple encoded data values can be packed into the processor register at the same time. All compressed September, October, November, and December data values for a single comparison can be taken concurrently (if they are in the same processor register), so predicate processing is much faster.

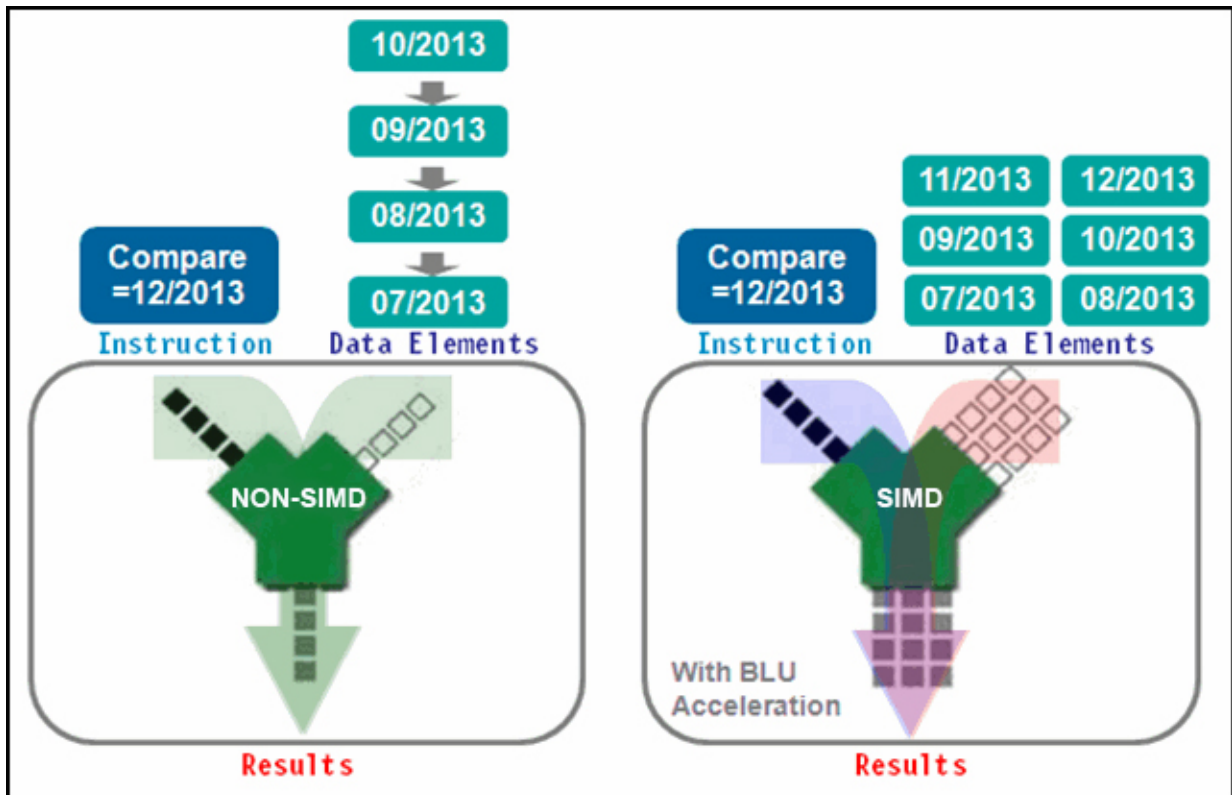


Figure 2. Efficient use of SIMD processor packing in BLU Acceleration

If SIMD-enabled processors are not available in your current hardware, BLU Acceleration still functions and speeds query processing using other features. In some cases, BLU Acceleration even uses its advanced algorithms to emulate SIMD parallelism to deliver similar results.

- Core-friendly parallelism

BLU Acceleration is a dynamic in-memory technology that makes efficient use of all processor cores in the system, so queries are processed using multi-core parallelism and scale across processor sockets. The idea is to maximize processing from processor caches while minimizing latencies.

In developing BLU Acceleration, much effort was invested in comprehensive algorithms for what is called *core-friendly* parallelism. The algorithms are designed so that data that is likely to be revisited is carefully placed and aligned in the processor cache lines. This maximizes the hit rate to the processor caches and increases the effectiveness of the cache lines.

Figure 3 shows an example. When a query is run, BLU Acceleration uses a separate agent for each available processor core to fetch individual column data values. Each agent can then work on different query functions. In the figure, each processor core (0, 1, 2, and 3) works on fetching data for different columns being queried. This is only one level of parallelism.

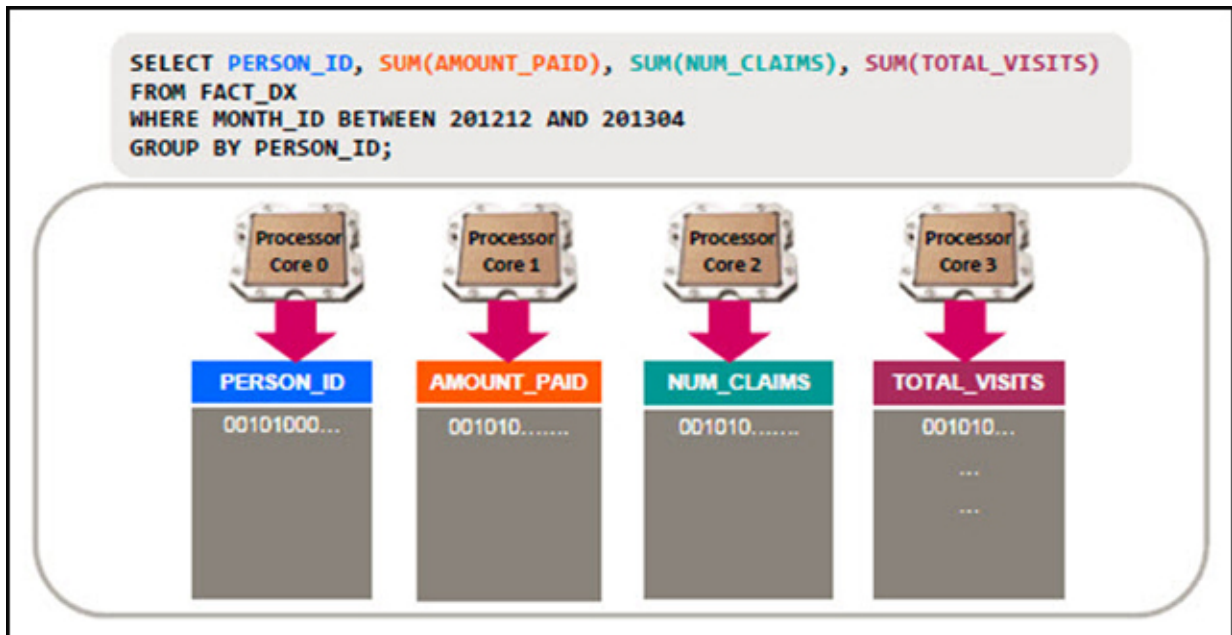


Figure 3. Multiple cores are used in parallel on multiple columns with BLU Acceleration

Figure 4 illustrates another level of parallelism. BLU Acceleration has intelligent algorithms that place and align data onto processor caches. This helps reduce the number of cache misses and maximizes processor cache hit rates. With data being available for reading from the processor cache without reading from the main memory or disks, query performance is improved significantly.

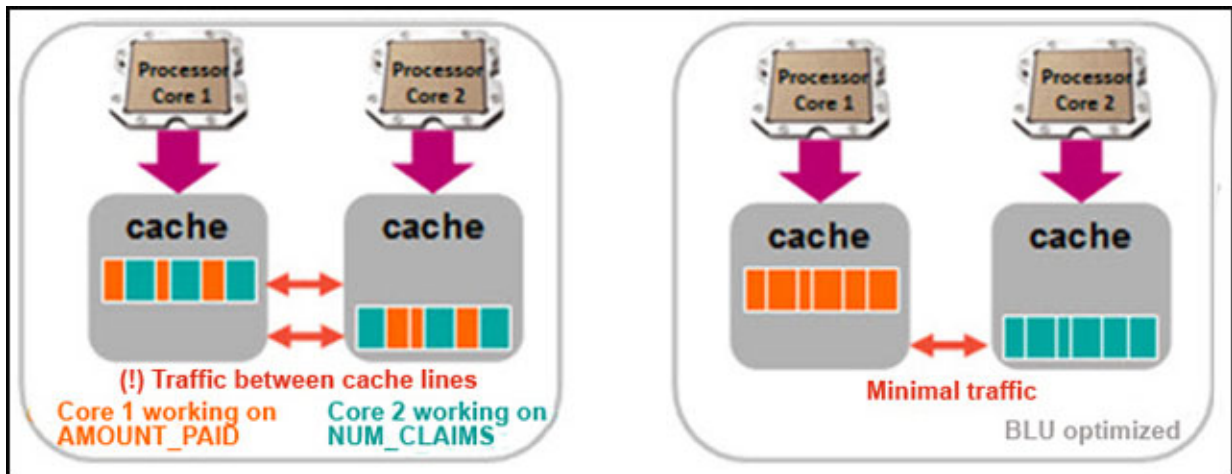


Figure 4. BLU Acceleration optimizes data alignment in cache lines

- Scan-friendly memory caching

BLU Acceleration is also designed for the demands of a big data world, where all of the target data for queries are less likely to fit into memory. Although DB2 can always benefit from having more hot (active) data in memory, it is not a requirement.

BLU Acceleration includes a set of big data-aware algorithms for cleaning out memory pools that are more advanced than the typical Least Recently Used (LRU) algorithms that are associated with traditional, row-organized database technologies. The new algorithms detect data patterns that are likely to be revisited and then hold those pages in the buffer pool as long as possible. When this is combined with BLU Acceleration's advanced columnar compression, even more hot data can fit into the buffer pool. Typical memory-to-disk ratios are 15 - 50 percent, but with BLU Acceleration, in many cases, up to 70 - 80 percent of active data can fit into memory.

- Data skipping

To speed query processing and reduce unnecessary I/O, BLU Acceleration can skip ranges of data that are not relevant to the current query. As an example, for a report about sales orders from December 2013, scanning them from October, November, and other months is unnecessary. BLU Acceleration goes straight to the sales orders for the target month and proceeds with the query.

This *data skipping* is accomplished by loading data into column-organized tables. BLU Acceleration tracks the minimum and maximum values in the various rows, using metadata objects called *synopsis tables*. When a query is run, BLU Acceleration looks at the synopsis tables for ranges of data that contain the value matching the query and skips directly to that point. Only necessary data is read or loaded into system memory, which speeds query execution by avoiding unnecessary scanning.

Solution architecture

The critical role of the data warehouse is to provide a central repository of structured data that is integrated from one or more sources. It enables a single version, or view, of the business, implemented as a single, scalable, and consolidated data warehouse. The data that is consolidated in the warehouse is available for analysis by all levels of the business. In addition, the warehouse contents can be subdivided into data marts that are specific to particular business units.

To support the workloads of today, the data warehouse architecture must be optimized for both traditional, deep analytic queries and for shorter, transaction-style queries. Figure 5 shows a high-level architecture of an analytics data mart that has been subdivided from a transactional database. Such data marts are created directly from the transactional database for use by lines of business. Because administration of BLU Acceleration is simple, data marts can be created quickly to enable rapid responses to new business requirements. DB2 with BLU Acceleration allows lines of business to run analytic reports in real time to improve their business decisions.

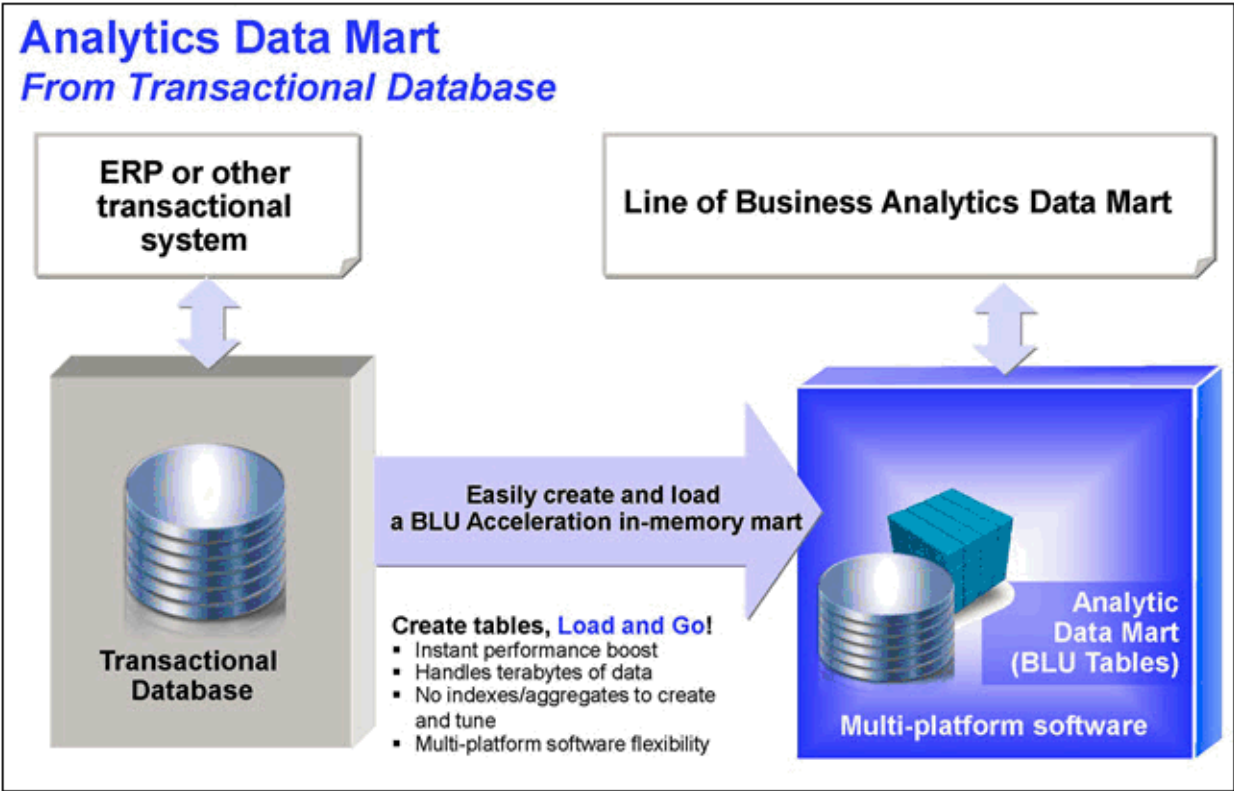


Figure 5. High-level architecture of an analytics data mart

DB2 with BLU Acceleration provides next-generation in-memory computing that meets rising demands of business users for fast analytics and immediate data reporting. Figure 6 shows a BLU Acceleration solution to offload analytic data marts from poorly performing, non DB2 data warehouses. This allows line-of-business users to use their existing OLAP application tools, such as IBM Cognos® Business Intelligence, and get faster reports in real time from the in-memory, column-organized data mart stores of DB2 with BLU Acceleration.

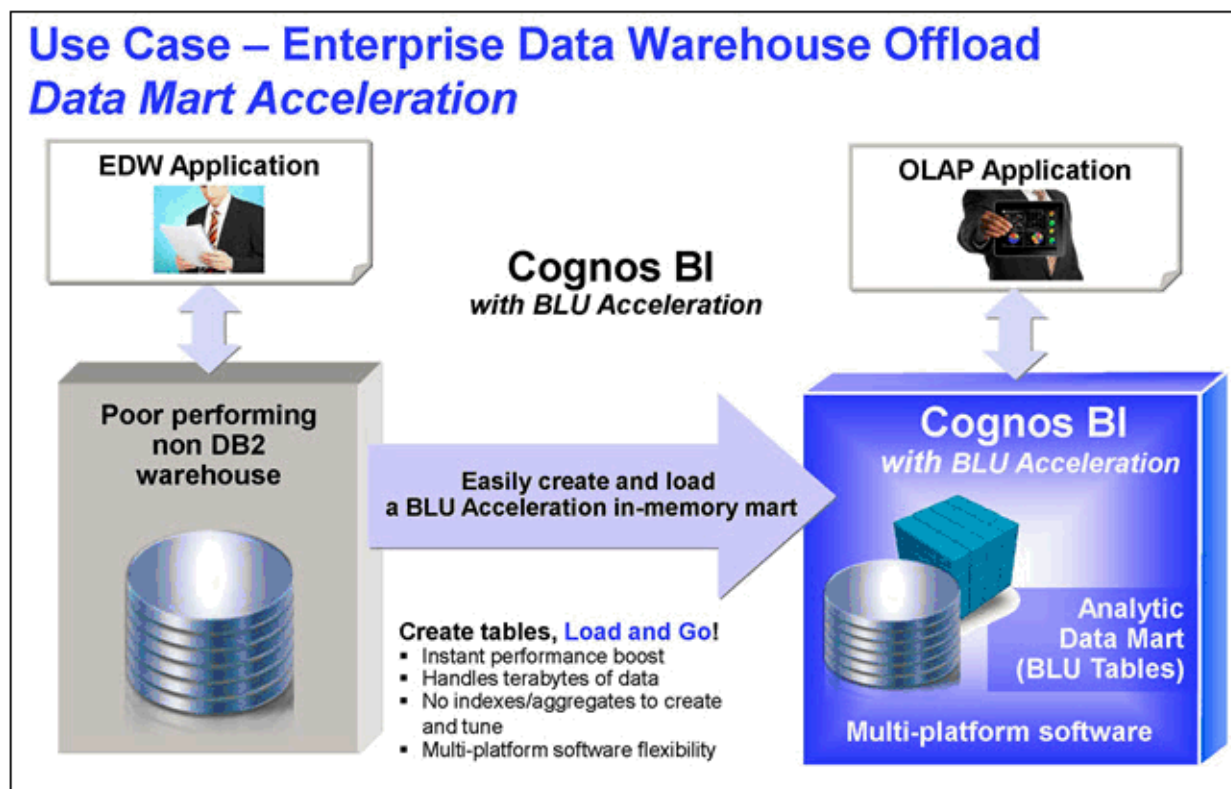


Figure 6. BLU Acceleration solution to off-load analytic data marts

Usage scenario

This example demonstrates how DB2 clients are observed to achieve significant performance and storage improvements using BLU Acceleration. The intent is to show what happens when a query is run on a large amount of data.

Assume that we have an analytics environment with the following attributes:

- Raw data: 12 TB
- Database server: SIMD-enabled processor, 30 cores, and 240 GB of memory
- Database: Two years of data (2012-2013)
- Sales order table: 100 columns

When we run a query to look for sales orders that shipped in December 2013, this is what occurs:

- Query
`SELECT SUM(ORDER_TOTAL) FROM SALES_FACT WHERE MONTH='201312'`
- Columnar Compression
From 12 TB of raw data, we conservatively assume a five times (5x) compression rate, reduces the size to 6 TB. This is a conservative assumption because BLU Acceleration is often observed with 10x compression rates.
- Column Store
From the total of 100 columns, the query accesses only those for ORDER_TOTAL and MONTH. This method of accessing column-organized tables reduces the target data to just two of the 100 columns, representing a mere 1/50th of the compressed data or 120 GB.
- Data Skipping
By using synopsis tables, BLU Acceleration skips to the data that matches the query predicate without decoding or evaluation processing. The query accesses just the ORDER_TOTAL for December 2013, or 1/24th of the records in the entire database, reducing target data to 5 GB (120 GB / 24 = 5 GB).
- Core-Friendly Parallelism
Data is processed in parallel across the 30 available processor cores, so each core needs to process only 167 MB of data (5 GB / 30 = 167 MB).
- Parallel Vector Processing
Data can be scanned using a SIMD-enabled processor to achieve faster performance. The amount of improvement depends on the processors being compared. Using a conservative assumption that our SIMD-enabled processor obtains each byte four times faster, DB2 will have to scan only about 42 MB of data (167 MB / 4 = approximately 42 MB).

To summarize, BLU Acceleration's performance and storage improvements make it possible for DB2 to have to process only 42 MB of the original 12 TB of data. This is roughly equal to a few high-quality digital pictures or a couple of social media cell phone apps. And with today's ultra-fast processors, the results of analyzing 42 MB of data can be returned almost instantly.

Integration

As with most IBM solutions, one of the beauties of DB2 with BLU Acceleration is the way it easily integrates with various complimentary solutions and products,

Cognos Business Intelligence

Cognos Business Intelligence and DB2 with BLU Acceleration work together. At the time of publication, all DB2 editions that support BLU Acceleration include five authorized user license entitlements to Cognos Business Intelligence 10.2.

Cognos Business Intelligence is a scalable enterprise business intelligence (BI) platform that empowers organizations with new levels analytic information to help drive better business decisions. It includes Dynamic Cubes technology that helps boost query performance by making more efficient use of memory. With Cognos Dynamic Cubes, aggregates, expressions, results, and data can all be stored in-memory. Ad-hoc queries or detailed drill-down reports can be completed almost instantly, even when run against terabytes of data.

Yet even though Cognos Business Intelligence provides techniques such as Dynamic Cubes to optimize analytics on large amounts of data, an almost unavoidable task is to completely bypass access to the underlying database, such as when the target data might not present in the Dynamic Cubes. The system then must access the required data from the relational database. DB2 with BLU Acceleration complements the Cognos optimizations with its dynamic in-memory column store technologies on the database layer. It is optimized for analytic workloads, providing an efficient data platform for accessing large amounts of data for specific columns when needed.

To learn more about using BLU Acceleration with Cognos BI, visit this website:
<http://www-03.ibm.com/software/products/en/bi-pattern-blu-acceleration>

IBM Power Systems

Using DB2 with BLU Acceleration on IBM Power Systems creates a combined hardware-software solution for faster in-memory processing. BLU Acceleration includes features specifically designed to use the capabilities of Power systems for big data and analytic applications.

Power Systems is an open, secure, and flexible platform that includes massive I/O bandwidth to deliver analytics in real-time. Built with the new IBM POWER8™ processor, the first processor designed for big data, Power Systems are further optimized to handle data processing needs like analytics and big data.

BLU Acceleration in-memory computing optimally leverages large cache sizes and memory bandwidth of Power Systems. This accelerates access to analytic results and reports for data-driven decision making. The BLU Acceleration Solution - Power Systems Edition further reduces deployment time and effort, offering preinstalled and pre-optimized server, storage, and software capabilities.

Learn more about the benefits of combining BLU Acceleration with Power Systems at this address:
<https://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&appname=gpatem&supplier=649&letternum=ENUSA14-0525>

SAP

DB2 for Linux, UNIX, and Windows (DB2) with BLU Acceleration can be integrated with SAP Business Warehouse (SAP BW) and SAP's near-line storage solution on DB2.

Creating column-organized database tables for the SAP BW objects provides several advantages:

- Reporting queries run faster when column-organized tables are used.
- Performance improvements can be achieved with little or no database tuning. Table statistics are collected automatically, and reorganizing tables or creating indexes for specific queries, statistical views, or database hints is usually unnecessary.
- Thanks to the faster query performance, the need for aggregates can be reduced or potentially eliminated. The SAP BW Accelerator might become obsolete, which greatly improves the extract, transform, and load (ETL) processing time.
- Except for primary key and unique constraints, column-organized tables have no indexes, so during ETL processing, hardly any time must be spent on index maintenance and index logging.

More information about using BLU Acceleration with SAP BW is at this website:
https://help.sap.com/docs/NO_PRODUCT/4c49a344277943ad91358094fdaf9765/d4fc57529daf785fe100000a441470.html

Supported platforms

The BLU Acceleration feature in DB2 10.5 is supported on the IBM AIX® and Linux x86 (Intel and AMD) platforms using 64-bit hardware. It shares the same minimum operating system requirements as DB2 10.5. See this location:

<http://www.ibm.com/support/docview.wss?uid=swg27038033#105>

Hardware information is in Table 1. Generally, the supported processors in DB2 10.5 are also supported for BLU Acceleration. The processors listed in the table include specific hardware optimizations from which BLU Acceleration can benefit. In the table, RHEL is Red Hat Enterprise Linux, and SLES is SUSE Linux Enterprise Server.

Table 1. Hardware and platforms suggested for DB2 with BLU Acceleration

Operating system	Minimum version required	Version to use	Suggested hardware
AIX	AIX 6.1 TL7 SP6 AIX 7.1 TL1 SP6	AIX 7.1 TL2 SP1 or later	IBM POWER7® or later
Linux x86 64-bit	RHEL 6 SLES 10 SP4 SLES 11 SP2	RHEL 6.3 or later SLES 11 SP2 or later	Intel Nehalem (or equivalent) or later

Ordering information

BLU Acceleration is available with DB2 10.5 Advanced Enterprise Server Edition, Advanced Workgroup Server Edition, and Developer Edition. To order, contact the Americas Call Centers, your local IBM representative, or your IBM Business Partner. To identify your local IBM representative or IBM Business Partner, call 800-IBM-4YOU (426-4968).

More information about ordering DB2 with BLU Acceleration is available at this address:

<http://ibm.co/1mkrh87>

In addition, potential customers can try DB2 with BLU Acceleration in several ways:

- *IBM DB2 BLU Acceleration Kit for Trial:* Install a fresh operating system and DB2 software stack together, with minimal system setup hassle. For more information and to download the kit, visit our Information Management Solution Portal:
<http://ibm.co/15zDz3C>
- *IBM DB2 with BLU Acceleration for Cloud Beta:* Access DB2 with BLU Acceleration on the cloud with your browser for up to 5 hours. To learn more, visit this website:
https://help.sap.com/docs/r/77e3a615f633437e9c4549891bb6bad7/CURRENT_VERSION/en-US/4a799891bf0d42c682ba890e1f80a2fe.html
- *IBM DB2 BLU Boot Camp or IBM Education:* Attend boot camps and education classes with presentations and hands-on lab exercises to learn about BLU Acceleration. Contact your IBM representative for details.

Related information

For more information about DB2 with BLU Acceleration, see these IBM Redbooks publications:

- *Leveraging DB2 10 for High Performance of Your Data Warehouse*, SG24-8157:
<http://www.redbooks.ibm.com/abstracts/sg248157.html?Open>
- *IBM Cognos Dynamic Cubes*, SG24-8064:
<http://www.redbooks.ibm.com/abstracts/sg248064.html>
- *DB2 Workload Manager for Linux, UNIX, and Windows*, SG24-7524:
<http://www.redbooks.ibm.com/abstracts/sg247524.html>
- *InfoSphere Warehouse: A Robust Infrastructure for Business Intelligence*, SG24-7813:
<http://www.redbooks.ibm.com/abstracts/sg247813.html?Open>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises is entirely coincidental.

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© Copyright International Business Machines Corporation 2014-2023. All rights reserved.

This document was created or updated on March 14, 2023.

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:
redbooks@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/tips1204.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:

AIX®
Cognos®
DB2®
IBM®
InfoSphere®
POWER®
Power Systems™
POWER7®
POWER8™
Redbooks®
Redbooks (logo)®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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.