



# Borrowing from Science to Revolutionize Business Analytics

Intel and IBM collaboration delivers a quantum leap forward with 25x faster performance in IBM DB2 with BLU Acceleration

Research engineers at Intel and IBM have a long history of working together to offer industry-leading technologies—many of them designed for the fastest computers on the planet, like the supercomputing clusters used for astrophysics or oil and gas exploration. Intel processors and software expertise, combined with IBM software designed to make the most of Intel technology, have enabled knowledge to be extracted from vast amounts of data. Now Intel and IBM have taken insight gained from these projects and applied it to challenges that other businesses in virtually every industry face on a daily basis.

For business, the era of big data is here. Companies are dealing with more data every day, and that data is arriving faster, from more diverse sources. To make sense of all this information at the speed of business and respond to new opportunities, companies need lightning-fast data access and analysis. But it's tough for traditional data warehouses to do the job.

## Accelerating results with columnar data store

Borrowing from the world of supercomputers and building on rich history of joint innovation, Intel and IBM engineers have collaborated to help translate torrential flows of information into business insights. The result is an IBM big data solution—optimized for Intel® Xeon® processors and developed with help from Intel engineers—that lets businesses run faster queries and gain more insight from their data.

DB2 10.5 features BLU Acceleration, which uses a revolutionary combination of columnar data store, memory optimization, and hardware exploitation to deliver a quantum leap in performance for analytics workloads. Online transaction processing, which acts on data for one customer at a time, is most straightforward when the data is organized in rows. But for analytical queries that need to scan through large sets of data, a column-based approach is faster. DB2 10.5 delivers both in one solution, so business users can get the rapid results they need. Analytics users can even ask a business question and draw on both row and columnar data at the same time to get the answer.

## Capitalizing on extreme compression

The experience of Intel and IBM engineers with big scientific and research databases also confirms that columnar data is much easier to compress. DB2 with BLU Acceleration uses advanced encoding for extreme compression. That translates to faster analytics, because it means more data can fit into main memory on the database server, where the server's processors can access it quickly. Of course, taking full advantage of this capability depends on having the right processor.

## Tapping into Intel processor advancements

As a result of Intel and IBM collaboration, Intel Xeon processors work together with BLU Acceleration to deliver the full benefits of



the technology and add even more speed to business analytics. The DB2 database engine can use the Single Instruction Multiple Data (SIMD) instruction set in the processor and register-friendly encoding from IBM to pack multiple data elements into a single register. The processor can then act on all of those data elements with just one instruction and get a lot more work done. For example, scanning four data elements with one instruction means the work can go four times faster.

DB2 with BLU Acceleration is also designed for improved parallelization, so the multiple cores within an Intel Xeon processor can divide up an analytic task into multiple threads and work on them simultaneously. And with the latest Intel Xeon processors, each processor can handle as many as 20 threads at once. It all adds up to faster answers to business questions.

Reliability, availability, and serviceability (RAS) features built into Intel Xeon processors help ensure the database server is highly available as well as high performing, and provide advanced data integrity for DB2 with BLU Acceleration databases. Intel Xeon processors offer a set of RAS features integrated in the processor silicon to provide error detection, correction, containment, and recovery in all processors and memory.

## Keeping performance high as the business grows

As businesses and their use of analytics expand, the scalable performance, memory, and I/O capacity of the Intel Xeon processors enable the business to adapt easily to changes in short-term demands and address requirements for longer-term growth. For example, Intel microarchitecture allows scaling

up to 10 cores per server socket and 30 MB of shared cache. Intel® Scalable Memory Interconnect (Intel SMI) and scalable memory buffers accommodate high-capacity bandwidth to help ensure that speed doesn't suffer as the size of the database grows. And Intel's advanced encryption instruction set accelerates encryption and decryption speeds, allowing businesses to extend their use of encryption to improve data security without sacrificing performance.

## Accelerating analytic results: An example

The combination of Intel Xeon processors and DB2 with BLU Acceleration can dramatically reduce the amount of processing needed to answer business questions. Here's how.

*The task:* Start with a data warehouse table containing 10 years of data, 2002 to 2011. Examine just the data from 2010, using a server with 32 processor cores.

*Size of the table:* 4 TB

*Effects of Intel and IBM technologies:*

- Compression reduces data size to 400 GB.
- Columnar access reduces the data to be searched to just 4 GB.
- Automatic mapping eliminates the pages without 2010 data, leaving 400 MB.
- Core-friendly parallelism means each processor core only has to scan 12.5 MB.
- Single Instruction Multiple Data (SIMD) and other technologies speed processing by a factor of four.

*The result:* Each core processes the data as quickly as if the table held only 3 MB of data, not 4 TB.

**"Intel is excited to see a 25x improvement in query processing performance using DB2 10.5 with BLU acceleration over DB2 10.1. To achieve these amazing gains, IBM has taken advantage of the Advanced Vector Extensions (Intel AVX) instruction set on Intel® Xeon® processor E5-based systems. Customers running this hardware can now immediately realize dramatically greater performance boost at lower cost per query."**

**—Pauline Nist, Intel General Manager,  
Enterprise Software Alliances, Datacenter  
and Connected Systems Group**

## Making smart use of data

Making the most of business data has become critical to the success of organizations across industries, and the collaborative relationship between Intel and IBM researchers is helping to meet the challenge. Designed to help businesses analyze rapidly growing volumes of data cost-effectively and extract more value in the form of action-oriented business insights, the newest generation of Intel Xeon processors and IBM DB2 with BLU Acceleration can deliver a dramatic increase in speed, exceptional simplicity, and operational cost savings for rapid return on investment.