

## ANALYTICS—ADAPTIVE, ACCELERATED, AND AVAILABLE

### *REAL VALUE DELIVERED THROUGH IBM BUSINESS ANALYTICS*

June 2014

A White Paper by

Dr. Barry Devlin, 9sight Consulting  
[barry@9sight.com](mailto:barry@9sight.com)

*By describing seven real customer implementations across four industries, this paper steps beyond marketing hype around analytics and provides compelling evidence of what it can deliver in actual implementations. To differing degrees, these implementations show triple-A characteristics required of analytics today to create business value: adaptive decisions and accelerated action based on all available information delivered via the optimal technology.*

*The example use cases, all built on IBM DB2 Analytics Accelerator, describe the various areas of business value delivered, from significant increases in data volumes handled to dramatic improvements in query response times, all contributing to business goals of faster and better decisions, improved customer support and lower costs, both in business areas and in IT support.*

*These cases offer nine insights into the path that organizations must take to progress from today's business intelligence to tomorrow's new, more powerful vision of analytics. This vision emphasizes the need to balance traditional BI characteristics such as flexibility of use and scalability with classical operational characteristics like reliability, availability and security when building analytic environments.*

## CONTENTS

- 2 Executive Summary
- 4 What are some real businesses doing?
- 7 Nine insights into triple-A analytics
- 9 Conclusions

Sponsored by:  
International Business Machines  
[www.ibm.com](http://www.ibm.com)



---

## EXECUTIVE SUMMARY

**W**hat to do now? Where to next? For today's enterprise, whether public or private sector, these urgent questions address immediate action and future expectation. The old demands of how and why goals are being met remain important, of course. Whether these goals are positive health outcomes, on-time arrivals, citizens well-served, or financial profitability, tracking progress has always been fundamental to good management. But the focus has moved on: from watching the rearview mirror to driving in the moment and anticipating the road ahead.

In popular conversation, this shift in attention from past to present and on to future is caught in a single word: *analytics*. But, what is it, exactly? Is it statistical analysis performed by data scientists? A new name for business intelligence, or souped-up data mining, perhaps? Is it about the current (operational) activity, some future (predictive) outcome, or both? How does it relate to reporting and traditional querying? Can a business person do it alone or does she need an army of programmers? Most importantly, what can it actually deliver, in terms of real business and organizational goals?

Unfortunately, like many trending and trendy words, a blizzard of marketing verbiage obscures the truth. The term can be—and usually is—applied in such a wide range of business and technical contexts that its meaning varies endlessly and continuously. But, let's take a different approach. Knowing the interconnected, ultra-complex, ever-changing, customer- and event-driven world of modern business, public or private, let's boldly define what we really need analytics to be:

***Adaptive decision making and accelerated action  
taking in every business activity, based on all available  
information, delivered via the optimal technology.***

**Triple-A  
Analytics**

Such triple-A analytics may sound like a tall order to the embattled CIO with a plethora of systems and a backlog of requests. However, the answer is hiding in plain sight: a dramatic simplification and consolidation of solutions using the exponentially growing power of technology to automate and accelerate every aspect of provisioning, securing and using data. Those familiar with mainframe technology will recognize this always-on approach. Those struggling with the management of the diverse and distributed hardware and software environment associated with current "analytics" would do well to take a closer look.

A key strength of the System z mainframe environment has long been its role as a central command and control point for all data and processing, the integrity of which is increasingly important to today's always-on business. Traditional distributed systems, on the other hand, often emphasize reduced commodity and open source component costs that are important in initial implementation and system expansion. However, today's price performant mainframe design actually incorporates such hardware and software components. It places their value and power "behind" a single, secure command and control point, where data provisioning and management is increasingly automated and accelerated. These characteristics are vital for the modern vision of triple-A analytics.

The introduction of the IBM DB2 Analytics Accelerator for z/OS (Accelerator) brought this revolutionary change, enabling a hybrid environment that was capable of both transactional and analytical processing on a single platform. The Accelerator complements DB2 for z/OS with a cost-effective high-speed query engine to run complex business analytics workloads. Operational processing in DB2 for z/OS is thus seamlessly combined with near real-time analysis and reporting using the Accelerator.

In brief, the DB2 Analytics Accelerator is a high-performance appliance that integrates zEnterprise and IBM PureData System for Analytics, powered by Netezza technology. This hybrid infrastructure blends the best attributes of symmetric multiprocessing (SMP) leveraging DB2 for z/OS with the best attributes of the hardware-accelerated massively parallel processing (MPP) architecture delivered by Netezza technology. Data originating in the DB2 for z/OS environment is automatically replicated from the mainframe to the Accelerator. The DB2 optimizer decides which queries to direct to the Accelerator for hardware-accelerated parallel query processing, while those that would not benefit from massively parallel processing continue to run within the highly-efficient DB2 for z/OS environment. The Accelerator is thus essentially transparent to applications and reporting tools querying DB2 for z/OS. Further details can be found in a previous white paper<sup>1</sup>.

*The IBM DB2 Analytics Accelerator enables a hybrid platform, joining the reliability and security of System z with the analytic muscle of IBM Netezza technology.*

The success of this marriage of the more traditional relational environment—where substantial quantities of business data already reside—to the new Accelerator is to be seen in its offspring: a series of innovative and productive use cases that take existing and often underutilized data and drive new value from it through faster, broader and deeper analyses. Customer and transaction data have long been at the heart of business intelligence, but often in periodic reporting and after-the-event problem analysis. DB2 Analytic Accelerator provides near real-time access to such data—often called operational analytics—that gives web sites and call centers the ability to respond and react immediately and relevantly to customers’ needs. Speed of response is also important for business analysts and managers who don’t have time to wait for reports that take ever longer as data volumes expand. Using ever larger and more varied data sets enables businesses to project trends into the future with more confidence and foretell individual customer behaviors—known as predictive analytics.

And the ability to do all of this in a single, integrated and secure environment offers benefits and cost savings to IT too—in reduced maintenance and in the use of existing queries from the DB2 for z/OS environment with minimal or no modification. Combining existing, secure data with data from external big data sources, seldom renowned for their dependability, in a reliable environment offers a level of control and governance that adds the seldom-mentioned fourth v—veracity—to the volume, velocity and variety normally listed.

How this approach works in practice is the theme of the following section. The golden thread through all the examples is *information-driven*: taking advantage of available data in the business to improve adaptive decision making and accelerate action taking. As queries return results in significantly shorter time periods, sometimes called “speed of thought”, business users move to more exploratory and iterative behavior that is ultimately highly productive but, more importantly, deeply innovative. Faster query turnaround times allow users to explore more data, dig more deeply, and acquire more useful and usable insights into past, current and likely future trends, as well as what to do about them.

*Information-driven: taking advantage of available data in the business to improve adaptive decision making and accelerate action taking.*

---

## WHAT ARE SOME REAL BUSINESSES DOING?

### INSURANCE—FUTURE HISTORICAL

Since the earliest days of data warehousing, insurance companies have been amongst its most eager adopters due to the central role that historical policy data plays in their business model. They also face some of the toughest challenges. Historical policy data is big data in terms of volume and variety. Velocity is now becoming a key business driver. Policies are becoming increasingly complex, demanding the storage of, and access to, ever more data. The needs of business users grow correspondingly. For sales to respond quickly, they require fast responses to complex quotations every day. This also requires analysts to run ever more intricate queries and reports as transaction volumes grow.

*A global reinsurance and insurance provider* needed a consolidated view of all open and finalized claims to set premiums and make accurate profit and loss calculations, based on a history of some 2.5 billion financial transactions that it had stored over the years. Running reports against 40 TB of data had previously required days. With the DB2 Analytics Accelerator, business users saw query results that had taken hours to run come back in minutes, with some results being delivered as much as 70% faster. That's the difference between waiting until tomorrow and making a decision today, delivering a significant competitive advantage to this business. Of course, end user satisfaction also increased dramatically. On the IT side, consolidating data on to a single platform has reduced management costs, as long-standing DB2 skills are applied to both operational and informational needs. Another benefit is the reduced mainframe CPU load as analytical queries are automatically directed to the Accelerator.

*A fast-growing European insurer* was determined to minimize IT costs and maximize profits while meeting transaction turn-around time and improving analytic support. With the Accelerator installed, they reported that queries that used to take five hours in the traditional environment now run in 20 seconds, and fully half of their queries now run 100 times faster. One specific example shows how a business process can be redefined. Previously, a particular sales report took a week to generate and could only be distributed late in the month. With the Accelerator, this critical report can be produced in a single day and distributed more frequently. This allows agents to proactively base behavior on ongoing trends to drive towards sales targets, rather than having to take reactive measures in the last week of the month. Furthermore, IT is empowered to react to new information requests from the business. Almost instantaneous access to business information drives data quality improvements compared to prior "fire and forget" report generation.

With such substantial improvements in making new insights accessible, the Business Intelligence (BI) team now has the ability to deliver process improvements that were previously unthinkable. Merging historical and operational data offers agents additional insight into how to structure a business response/quote in near real time. More extensive fraud detection becomes economically viable by exploiting externally sourced social media or public data. The Accelerator offers the power to integrate the analysis of the past, present and future for immediate business decisions. And the benefits are not confined to the analytic environment: with analytics now running in the Accelerator, operational transaction response times have been reduced by 20%. As transaction volumes increase, this insurer is very pleased to have deferred investment in additional mainframe resources.

*At one US health insurer*, the improved businesses efficiencies and results from the Accelerator have been so dramatic that users have started to talk about the "Magic Box". As business demands increased on the existing system, many reports took several hours or more to complete and used signif-

ificant computing resources. With the Accelerator, elapsed time for one query was reduced from over 24 hours to 72 seconds; for another, the improvement was even greater at 4800 times. Because no changes were required to existing application programs, the DB2 Analytics Accelerator could be integrated seamlessly and effortlessly into the existing business processes.

One analyst eliminated hours of repetitive work on one regular report against one particular large table. Previously, he had to split his queries across multiple data subsets, each small enough to return results in under an hour. He then exported, cut and pasted the results into a single report. With DB2 Analytics Accelerator, he could run a single query returning over 7 million records in less than 15 minutes. Some queries that previously timed out, wasting valuable computing and analyst time, or were considered too costly to run at all could now return valuable business results in minutes. The opportunity to gain new insight into operations and recognize and implement optimizations was significant.

## **BANKING ON CUSTOMER DATA**

Like many other financial institutions, banks have entrusted their transaction data to zEnterprise for reliability and security of their financial systems. Worldwide, 92 of the top 100 banks use System z and over 70% of financial customer and transactional data originates in this environment. Creating an analytic solution within the zEnterprise environment is an attractive option, avoiding the need to copy data to a different environment that requires new skills and increases complexity. Data warehouses have a long history on DB2 for z/OS; its predecessor on MVS hosted some of the earliest warehouses ever built. However, there were limitations. Today, analytics have become business critical and protecting the access to and proliferation of the data is vital. Retaining data within the zEnterprise environment makes more sense than ever before.

*A major European banking group* has recently been able to develop and deliver a novel online analytics solution, thanks to the Accelerator. Transaction and customer data for the group's largely independent constituent banks resides in DB2 for zOS on the mainframe and is accessed for all operations from the group's near-2,000 branches. However, data volumes have long dictated that only current day data is held online. At the end of every day, the day's data was archived to tape and data more than 2 months old was deleted from the database. Analytical access to and, even, operational use of such archived data could only be offered by reloading data from tape into the online environment. This procedure introduced a day's delay into any business need for older data, with significant productivity and competitive effects. The Accelerator simply replaced the entire tape archive and retrieval process with an online solution. The change was completely transparent to the users—except, of course, for the dramatic improvement in turnaround time for queries. Not only did this approach remove the full-day delay for data retrieval, but it also significantly accelerated the running queries and reduced the DB2 storage and CPU load on the System z. A further important benefit was the fact that existing queries and reports could be used without modification of previous DB2 for z/OS SQL.

Much of what is written about big data these days focuses on the use of social media and other novel data types to understand customer sentiment and predict possible behavior. What is often overlooked in the excitement is that a key—actually mandatory—first step is to move the organization to a data-driven mindset based on existing internal data.

*Another large European bank* started their big data project to get the maximum value from their existing and extensive amounts of customer data which was housed on System z. Such data is, of course, highly sensitive, so cast-iron security is a strong requirement. By retaining this data within the secure vault created by the System z environment, data that is maintained in the Accelerator inherits DB2

for z/OS data attributes including security and reliability. Thousands of business users can thus benefit from this platform's top class security features, as well as the enablement of high-speed analysis of real-time data. Users from both customer-facing and product development business areas utilize this data to develop new insights and improve performance and timeliness. For example, the Accelerator has enabled the bank to reduce the time taken to introduce new marketing campaigns from months to weeks. Next steps will focus on delivering deeper analytics across all areas of the business, including using IBM predictive technology in SPSS to uncover emerging trends in customer behavior, using long-term data horizons to estimate probabilities of default as required by the Basel III<sup>2</sup> framework, as well as analyzing and improving the effectiveness of marketing across all contact channels.

## **RETAIL—RIGHT PRODUCT, RIGHT PLACE, RIGHT TIME**

In retail, interest is often focused on the dominant megastores, with their world-spanning supply chains and their predictive analytics down to the last cereal packet on the shelves. But, smaller stores, especially those which are part of co-operatives or chains, also demonstrate the value of getting the product mix and volumes right for varied and changing customer needs. The result is the profitable sweet point of treading the thin lines between disappointing gaps on shelves and bulging stockrooms of unwanted merchandise. In addition, effective cross- and up-selling at the point of sale becomes practicable and profitable.

*A national fuel supplier* of oil and other energy products also offers a broad range of automotive goods and services and a wide selection of household items, food products and other merchandise through its approximately 570 service stations. A loyalty card program had enabled the collection of large quantities of detailed customer purchases and product sales across both channels. The challenge was to analyze this data efficiently and provide the required information to store cashiers that would enable them to suggest additional relevant purchases to both loyalty and walk-in customers within the few tens of seconds that a typical transaction takes. Beyond velocity, volumes were also important: this chain serves a quarter of a million customers per day and its transaction database is in excess of 20TB. With the Accelerator in place, cashiers are instantly shown complementary products to those in the customer basket, resulting in additional revenue through cross-sell and up-sell opportunities. For loyalty card holders, previous purchase history is displayed, allowing cashiers to improve their judgment on what to cross-sell. Furthermore, using technology within IBM Cognos Business Intelligence and IBM SPSS Predictive Analytics software enables deeper analysis of sales data, so purchasing managers can track and predict product- and brand-level sales trends across regions and stores to optimize stock levels and deliveries.

## **UTILITIES LIKE TO CONSERVE POWER TOO**

*An electric utility company* with more than five million customers in the United States took advantage of a necessary hardware upgrade from two older IBM z10 servers to zEC12s to place new focus on operational analytics. The addition of the DB2 Analytics Accelerator provided the opportunity to integrate business insights with operational processes and also to accelerate the performance of transactional workloads. While the mainframe upgrade gives the company improved overall performance and efficiency, the Accelerator makes a significant contribution as well. Analytics often demands significant computing power, often at unexpected times or under unpredictable circumstances. Transparently offloading such processing to the Accelerator both reduces and evens out the operational processing load on the mainframe, reducing CPU contention for development and test, as well as other lower priority workloads. The IT department also reduced database maintenance because no indexes needed to be created nor reorganizations run on the analytical data. The result is an all-round performance improvement for all types of workload.

---

## NINE INSIGHTS INTO TRIPLE-A ANALYTICS

The preceding use cases illustrate how new technology can deliver real benefits to an organization by enhancing existing processes or transforming the value a business derives from its data. A particular game changer for an organization occurs when the speed at which insights are derived is dramatically increased, significantly enhancing the value of data already stored. While there is no single, one-size-fits-all set of drivers for accelerated analytics, we can observe a general progression on a path from basic functionality to empowered and extensive usage.

1. Traditional BI reporting and querying is a vital and often early component of triple-A analytics implementations, providing a baseline set of information for more advanced use cases. Faster and more efficient BI alone can justify some projects.
2. Triple-A analytics can also be initially driven by broader planning needs, such as scheduling production or improving stocking levels, refining marketing campaigns, and so on. Such demands move the focus from past performance to future possibilities and enhance the value of analytics.
3. Optimized hardware and software within the mainframe environment enables implementation of analytics across a wide range of business sizes and complexities at an affordable cost. Operational cost savings in the traditional mainframe technology can fully offset the costs of implementing analytics in the optimized technology.
4. Triple-A analytics is applicable to all industries, including government and non-profit, because every organization needs to optimize its business and activities. Irrespective of the industry, adaptive decisions and accelerated action are central to development, product sales, and customer interactions and are key to understanding individual transactions and emerging trends.
5. The substantial improvements in query performance at the heart of triple-A analytics enable two broad usage types: (i) run existing queries faster/cheaper and (ii) run existing or new queries on volumes of data previously impossible. Broadly speaking, the first type improves business performance in existing operations; the second enables and drives new business processes.
6. Faster and easier access to data offers new opportunities to target real-time business needs. Often termed Operational BI or Decision Management, these needs require a combined operational-informational environment, implementation of operational processes and service level agreements. These aspects are well understood and widely supported in the integrated, System z world. The DB2 Analytics Accelerator is thus an ideal extension to enable these hybrid workloads.
7. The ability to quickly and easily run previously slow, challenging or impossible queries is not only about productivity; it can energize business analysts and other users to new levels of creativity. Seeing results that took several hours to generate returned in minutes, or even seconds, changes the psychological processes involved in problem solving, speeding up insights and driving iterative exploration of the data rather than report delivery. Users can envisage new ways of combining and using data, and they can ask question after question, allowing process innovations to emerge.
8. For manufacturers and utilities, triple-A analytics offers the ability to track performance, optimize production, enhance maintenance, and generally improve the entire supply chain. Internal machine data is growing in volume and value as sensors become more sophisticated. A wide range of

*Triple-A analytics is applicable to all industries, including government and non-profit.*

devices, from simple electricity meters to sophisticated jet engines, is already delivering data at high speed from the Internet of Things. With data arriving at such velocity, real-time analysis and action is the first priority. Statistical analysis offers the opportunity for continuous improvement and the ability to detect emerging problems before they become serious.

9. Customers, products and transactions are the core data for all analytics. This is augmented by so-called “big data”—human-sourced information<sup>3</sup>, such as social media for trends in sentiment, and machine-generated data from mobile devices for location. Taken together, these three types of data represent the final frontier of triple-A analytics. While basic preprocessing and analysis of big data is typically performed in Hadoop, further analysis often requires consolidation with the traditional data warehouse environment. With DB2 v11 and the Accelerator, business users can gain seamless access to all this data. In effect, the System z becomes the managed, secure information hub of the organization.

*The consolidation of traditional and new data types for combined analysis is the final frontier for triple-A analytics.*

From one point of view, proceeding down this list provides an insight into the ability of a business to take advantage of and drive significant value from an investment in triple-A analytics. Leading-edge businesses find opportunities in the last three areas mentioned. On the other hand, the list also describes an evolution in business needs and IT functionality that has been ongoing for almost three decades now and promises to continue for the foreseeable future. The time has arrived to consider how business and IT must cooperate to drive maximum benefit from the broad endeavor that analytics represents.

### ***FROM TRIPLE-A ANALYTICS TO DB2 ANALYTICS ACCELERATOR***

Business-critical operational applications demand high levels of security, reliability and availability and are traditionally run on mainframe and other large, centralized systems such as System z. Business intelligence was often transferred to distributed systems, for a combination of reasons, including cost issues and performance gains for particular types of query processing. This traditional, bimodal implementation approach must be reconsidered in the light of three important factors:

- Today’s business demands deeply integrated operational-informational environments that enable extreme agility and speed in meeting customer needs and in responding to market changes
- Analytical and BI needs are now considered mission-critical for all mid-sized and larger businesses across all industries, demanding the levels of reliability, availability, security and management traditionally reserved for transactional systems
- While big data growth is overwhelmingly in machine-generated data and human-sourced information (associated in the market with Hadoop), it is increasingly clear that full, high-value use of such data depends on integrating it with traditional process-mediated data

*Triple-A analytics offers businesses new ways of working in a highly-interconnected, data-rich world.*

Although some analysts and vendors propose that these factors lead to a solution where all data is brought into a Hadoop-based data lake (or reservoir), the totally undifferentiated environment, the immaturity of the ecosystem, and the lack of data management experience all suggest that such an approach cannot succeed. A more realistic approach is to blend existing systems which offer the required characteristics. This is the approach taken in the IBM DB2 Analytics Accelerator, which combines the operational strengths of the System z with Netezza-powered technology that is optimized for analytical tasks.



---

## CONCLUSIONS

In a rapidly evolving business environment, driven by ever more powerful technology, pervasive networking and an explosion of information, traditional business intelligence has become table stakes to compete for business or deliver acceptable service. Aspiring market leaders, whether in profit or non-profit endeavors, must strive for increasingly rapid and accurate insights to enable “speed of thought” action. While these directions undoubtedly require business process reinvention and deep organizational recreation, the fundamental starting point is adaptive, accelerated and highly available analytics.

As we’ve seen in the example use cases explored in this paper, businesses in a wide range of industries have already achieved extensive and significant improvements in query response times, volumes of data handled, and results delivered by using a modern, hybrid environment. Such improvements are the foundation for new analytic journeys on which they are embarking: from standard reports to situation-specific analyses, from waiting overnight for answers to enabling instant information exploration, from yesterday’s data to real-time information. In many cases, the change in behavior of the business analysts doing this work is substantial enough to consider changing their role to that of the proverbial data scientist! They are not playing with Hadoop; they are using their existing skills to new effect. They are breathing new life into old data. And they are doing it in a cost-neutral implementation, extending existing DB2 on z/OS systems with the DB2 Analytics Accelerator.

While offering a powerful environment in its own right, the IBM DB2 Analytics Accelerator provides clear insight into the data architecture that is demanded by the big data explosion and the coming tsunami of information from the Internet of Things. It is already becoming obvious that no one technology—be it relational or Hadoop—can meet the storage and processing demands of this challenge. What the Accelerator shows is how these different technologies must be blended into hybrid systems with a combined control and management supervisory function, in this case, DB2 for z/OS. And what we also see is the value of extending the reliability, availability and security of the mainframe to new analytic needs.

*IBM DB2 Analytics Accelerator shows the data architecture that is demanded by the big data explosion and the coming tsunami of information from the Internet of Things.*

Triple-A analytics emphasizes three key characteristics of future-proof analytics—*adaptive* decisions and *accelerated* action based on all *available* information delivered via the workload-optimized environment. The customers featured in this paper demonstrate that these characteristics already exist and are delivering real business value.

---

Dr. Barry Devlin is among the foremost authorities on business insight and one of the founders of data warehousing, having published the first architectural paper on the topic in 1988. With over 30 years of IT experience, including 20 years with IBM as a Distinguished Engineer, he is a widely respected analyst, consultant, lecturer and author of the seminal book, "Data Warehouse—from Architecture to Implementation" and numerous White Papers. His new book, "Business unIntelligence—Insight and Innovation Beyond Analytics and Big Data" (<http://bit.ly/BunI-Technics>) was published in October 2013.



Barry is founder and principal of 9sight Consulting. He specializes in the human, organizational and IT implications of deep business insight solutions that combine operational, informational and collaborative environments. A regular contributor to [BeveNETWORK](#), [TDWI](#) and other publications, Barry is based in Cape Town, South Africa and operates worldwide.

Brand and product names mentioned in this paper are trademarks or registered trademarks of IBM and other companies.

---

<sup>1</sup> Devlin, Barry, "Integrating Analytics into the Operational Fabric of Your Business", April 2012. <http://bit.ly/IDAA-WP>

<sup>2</sup> Basel III is a global regulatory standard on bank capital adequacy, stress testing and market liquidity risk.

<sup>3</sup> Devlin, Barry, "Business unIntelligence—Insight and Intuition Beyond Analytics and Big Data", October 2013, Technics Publications, New Jersey. <http://bit.ly/BunI-Technics>