IBM Analytics
White paper

Cross-industry

Make your most valuable data more valuable with IBM Real-Time Analytics

Enhance the value of analytic insight with live enterprise data

Contents

Live data for real-time insight

Enterprise data is one of the most valuable of an organization's assets. It is the data, created and accessed by operational applications, that chronicles an organization's activities. Frequently, that data is stored and managed in an IBM® z Systems® environment. To derive maximum benefit from this high-value data and to profit from the transformational insight it can provide, the data needs to be accessed where it lives, in real time.

Data analytics — the platform matters

People in the real estate business often remark that the value of a property comes down to three things: location, location and location. Data is no different. Understanding where your data is located and using it in place can be a huge differentiator, particularly as organizations look to increase competitive advantage with more real-time analytic insights.

For many organizations, the bulk of their most valuable enterprise data is managed on IBM z Systems because of the platform's unparalleled reputation for reliability, availability, scalability and security. Businesses are increasingly focused on finding new ways to leverage their enterprise data and gain even greater value from it. Real-time analytics presents an enormous opportunity to do just that.

What is real-time analytics?

Real-time analytics is insight delivered fast enough to make a difference.

This definition covers a wide range of possibilities. Consider, for instance, a person using an analytics tool to interact with a database containing years of sales history information about his organization's products. The analytics tool, by way of a graphical interface, generates a complex query, which is sent for execution to the database system.

This complex query might involve scanning millions of data records. If it takes hours for the query response to be returned, the user might have some trouble recalling exactly what led him to initiate the query in the first place.

What if the answer could be returned in less than a minute, with the question still fresh in the user's mind? That would be much more likely to result in useful insights, leading to further refinement and follow-on queries that ultimately create actions that achieve breakthrough results. That's one example of insight delivered fast enough to make a difference.

This scenario — turbocharged performance for ad hoc, decision-support analytics — is certainly attractive. However, too many organizations don't see beyond this application of real-time analytics technology. Such a limited view of real-time analytics could cause an enterprise to miss the benefits that can be realized through an even more transformational approach: that of infusing analytics into core business processes.

Enhanced in this way, those processes don't just run the business. They elevate the business beyond competitors, and extend it to new clients, new markets and new revenue streams. This is the promise of what we call in-transaction analytics — a subset of real-time analytics.

Real-time and right-time analytics

In-transaction analytics is essentially the convergence of real-time and right-time analytics. While real-time analytics relates to immediacy of response, right-time analytics refers to the most opportune time to analyze data relative to its creation.

There are, of course, situations in which analysis of data created years ago is very useful and appropriate — the previous example of a user analyzing sales history information would be in this category. But there are other situations in which the value provided by analytics increases as the time between data creation and data analysis decreases. In-transaction analytics can reduce the length of time between data creation and data analysis all the way to zero — within the very transaction in which data is created, data can be analyzed in real time. That, in turn, opens up powerful possibilities.

In-transaction analytics aims to use the power of data analysis in what is known as the transactional moment. That is the time during which a transactional event is still in progress. Here are a few examples of transactional events:

- Someone checks out their purchases from a store
- Someone is on the phone with a customer care representative
- An insurance claim is being adjudicated
- A credit card purchase is being processed
- An electronic payment is being processed

In-transaction analytics uses data analysis to achieve a more favorable outcome for a transactional event, even as the event is in progress. In the case of the person checking out purchases from a store, the more favorable outcome might be a larger purchase.

For a call being handled by a company's customer care department, the more favorable outcome might be a customer who feels more pleased to do business with the organization, as opposed to possibly severing the business relationship.

For the adjudication of an insurance claim, the more favorable outcome might be the detection, and subsequent avoidance of, what might have been an over-payment.

For the processing of a credit card purchase or an electronic payment, the more favorable outcome might be detection, and subsequent prevention of, what might have been a fraudulent transaction.

The goal of in-transaction analytics is to detect an undesirable occurrence in time to prevent it, or recognize an opportunity in time to act on it. In either case, data analysis is the mechanism of discovery of a potentially undesirable occurrence or of a potential — and perhaps fleeting — opportunity. In both cases, time is of the essence.

Why? Because to make a difference in the outcome, the insight gained through in-transaction analytics has to be available before the transactional event in question has completed. Some transactional events—a person shopping in a store or online, for example—can have a duration of several minutes. On the other hand, transactional events such as adjudication of an insurance claim or processing of a credit card purchase can be completed a fraction of a second after they begin.

For situations that demand an ultra-fast analytic response time, the only feasible solution for in-transaction analytics (especially if analytics are to be incorporated into the processing of all—versus a sampling—of an application's transactions) is often in-transaction scoring, which involves invoking a predictive model from within a transaction. In-transaction scoring delivers the ultimate in low-impact, in-transaction analytics in terms of effect on CPU and elapsed time.

Analyzing data where it lives

The IBM z Systems platform is an ideal choice for managing data generated by an organization's most mission-critical applications. From both a real-time (immediacy of response) and a right-time (latency of data analysis relative to data creation) perspective, analysis of z Systems-based data on the z Systems platform maximizes capabilities.

Besides delivering the ultimate in real-time and right-time analytics capabilities, analyzing z Systems-based data on-platform, as opposed to copying the data to another system for analysis, offers the following advantages:

- Improved timeliness of data being analyzed
- Enhanced security
- Reduced data replication costs (these can often be substantial)
- Improved data governance—analyzing z Systems-based data on-platform can mean fewer copies of the data. Fewer copies mean better control. Better control leads to greater trust in the data. (For more information on protecting data used in analytics read the white paper "Is your business risking a USD 100 million loss?")

IBM z Analytics technologies deliver real-time analytics capabilities

Multiple hardware and software offerings, from IBM and from other vendors, provide a wide range of high-performance, high-function, real-time analytics capabilities that can help organizations gain greater value from enterprise data on the z Systems platform. Here are six examples:

- The DB2® Analytics Accelerator (Accelerator) enables cost-effective, high-performance analytics on the platform on which z Systems based data lives. Unlike analytics appliances that introduce one more data island in an IT infrastructure, the Accelerator is an extension of IBM DB2 for z/OS®.

  First and foremost, the Accelerator helps provide outstanding response time for highly complex, data-intensive queries directed to the front-end DB2 system. (The front-end DB2 system automatically routes queries to the back-end Accelerator when it determines that doing so would provide the best performance for query execution.) Queries that might otherwise run for hours can be completed in seconds. When a query is routed to the Accelerator, almost all of the CPU cost associated with query execution is shifted to the Accelerator from the front-end DB2 system.

  A relatively new feature, known as Accelerator-Only Tables, can help improve the speed and reduce the cost of iterative analytics processing, which involves an initial query result set that becomes the target of another query. The result of that query is accessed by yet another query, and so on.

  A related offering, the DB2 Analytics Accelerator Loader, makes it easier to get data from outside DB2 for z/OS—from IBM IMS databases, from VSAM files, from other vendors' z/OS database systems, even from non-z Systems sources—into the Accelerator for analysis, providing a broader range of data from which powerful insights can be derived.

- IBM SPSS® software provides a predictive analytics capability that lends itself to in-transaction analytics scenarios. How? First, the data mining function of SPSS can be used to find patterns within a set of data records. (For example, when A, B, and C are true, D is likely to happen.)

  Next, SPSS can turn a detected data pattern into a predictive model. This predictive model can be used to score a transaction — signaling, for example, that a particular payment transaction is likely to be fraudulent in nature, or that a particular customer in a particular situation is likely to be amenable to buying a certain product when certain other products are set to be purchased.

- If a predictive model is to be used within a very short-duration transaction, such as the processing of a purchase made with a credit card, it must be invoked in a way that minimizes its impact on CPU and elapsed time. One way to do that is through technology offered by Zementis, an IBM partner. Using the Zementis technology, an SPSS predictive model can be invoked to score a transaction while adding perhaps a millisecond or less to the transaction's response time and CPU cost. That makes in-transaction scoring feasible for even very high-volume transactional applications with very stringent throughput requirements.

- A score received via invocation of a predictive model is only useful if it can be acted on. Acting on a score— programmatically altering course based on the value returned by the predictive model — requires new application functionality. That presents an opportunity to develop the new functionality more efficiently and effectively, versus traditional approaches, using the IBM Operational Decision Manager (ODM) for z/OS.

  Instead of hard-coding business rules and decisions into programs (an approach that works against rule consistency and hinders rule re-use), ODM enables teams of developers and business people to define rules and decisions, such as what to do in response to a score received in-transaction from a predictive model, and validate and test those rules and decisions.

  These rules and decisions can then be deployed to an engine through which they can be invoked by both transactional and batch programs. ODM also provides facilities for rule governance (including versioning) and for the monitoring of rule and decision invocation and execution.

- Linux for z Systems provides an opportunity to run analytics software — whether IBM products such as IBM Cognos® Analytics, DB2 QMF (Query Management Facility), and SPSS, or other vendors' offerings, or open source tools — closer to z/OS-based data sources. With that proximity comes better performance, aided by HiperSockets, which are memory-to-memory communications links between logical partitions (LPARs) within a z Systems server, such as a Linux LPAR with Cognos Analytics accessing DB2 data in an adjacent z/OS LPAR.

- Spark, from the Apache Software Foundation, is another open source technology that is part of the z Systems real-time analytics picture. Spark, which can run on a z Systems server under either Linux or the z/OS operating system, is a tool for working with data versus persisting data (data persistence is the role of database management systems).

Spark can ingest data from multiple different sources, including relational databases (such as DB2), Hadoop-managed data stores, and—via connectors available from IBM—from z/OS sources such as IMS databases, VSAM files, and ADABAS databases. This capability makes Spark attractive from a data integration perspective.

Performance is another key strength of Spark. It can build in-memory structures with data read from one or more sources, and subsequent access to these data structures can deliver very fast response times for queries. In particular, Spark can greatly speed the execution of queries targeting data obtained from a Hadoop data store, versus accessing that data via Hadoop's native MapReduce interface. Add to this Spark's machine learning features and its capacity for processing data streams and you have a powerful data analytics framework that can be applied to a wide variety of use cases.

Here are just a few of the ways in which z Analytics technology can be put to work to help organizations run better and faster.

Making effective use of z Analytics technology

The DB2 Analytics Accelerator is used today in a variety of ways by enterprises in different industry segments around the world. A number of companies have found value in using the Accelerator to augment an existing analytics infrastructure. In such cases, the Accelerator provides a resource that complements, for example, a data warehouse, by providing users with high-performance, ad hoc query access to raw (that is, non-transformed), atomic-level (that is, non-aggregated) operational data, without imposing a processing burden on a front-end z/OS system. An example of this approach is the use of an Accelerator to provide users with query access to DB2 data that supports a vendor-supplied ERP application.

Figure 1: Diagram showing the components of the z Systems architecture with the DB2 Analytics Accelerator

Some organizations have used the Accelerator to replace an existing analytics infrastructure that has become too expensive and too difficult to manage. Especially when z/OS-based data is copied to other platforms for analysis, costs can escalate in several areas such as data replication software, analytics platform hardware and software costs, network resources and system support personnel. Additional costs associated with off-platform analysis of z Systems-based data relate to security and governance—more copies of data, on more platforms, make it more difficult to protect data and to ensure its accuracy.

Data latency problems are another concern, as users increasingly want the data they use for analytics to be very current with respect to source data. Yet another issue associated with the move-and-analyze approach is the proliferation of data analytics islands that eliminate value obtainable through analysis involving queries that could span multiple domains in an integrated, z Systems-based database.

By using an Accelerator to enable analysis of enterprise data in-place, as opposed to moving z Systems-based data to separate analytics systems, organizations have lowered costs, improved analytics performance and enhanced analytics capabilities.

A number of companies have used the Accelerator to implement a query-able data archive, taking advantage of the Accelerator's High-Performance Storage Saver feature. This feature can be used with a DB2 for z/OS table that is partitioned on a time-period basis (for example, one week per partition) so that only the current partition (for example, this week's data) is physically present on both the front-end DB2 system and the back-end Accelerator.

All other partitions—and there could be hundreds or thousands of these for one table—are physically present only on the Accelerator. The result is high-performance, cost-effective query access to historical data that might have previously been stored in an offline archive. The attractive economics of this Accelerator use case owe much to the fact that the Accelerator's integrated disk storage tends to be significantly less expensive than the high-end disk subsystems typically used with z Systems servers.

An insurance industry real-time analytics use case

An insurance industry scenario can illustrate how multiple z Analytics technologies can be used together to enhance operational capabilities in a transformational way.

Consider the problem of insurance claim over-payment. Over-payment of claims can be tied to fraudulent activity, or to honest mistakes on the part of claims submitters, or to mistakes made by an insurer in adjudicating claims. Whatever the cause, over-payments are a costly problem.

Complex queries that analyze claim payment information can help to identify incidences of over-payment, and executing these queries on an Accelerator can greatly reduce run times, so that actions to recover an over-payment can commence sooner rather than later following claim payment processing.

Taking steps to recover claims over-payments sooner after they have occurred rather than later is a step in the right direction. Better still, however, would be detection and prevention of over-payments before they occur. This can be achieved by using SPSS to develop a predictive model that can be invoked within the claims processing application to generate a score for payment transactions. That score, which in many cases can be returned in a thousandth of a second or less, can identify transactions that would be highly likely to result in over-payment if processed the usual way.

Those transactions can then be diverted from the main payment processing stream and sent to a queue, from which they can be pulled and examined by claims adjudication experts to determine proper payment. Figuratively, this can be thought of as using the predictive model to create an off-ramp from a claims processing superhighway, with tagged claims (that is, those predicted to result in over-payment) directed to a parking lot for subsequent examination and resolution.

Preventing over-payments before they occur can help deliver major cost savings for the insurance company. But the situation can be improved even further by taking an additional step: that of building the response to a "probable over-payment" predictive score into the claims payment processing application itself, versus relying on manual examination of flagged claims by adjudication experts.

This autonomic capability can be developed with greater efficiency and enhanced consistency, compared to traditional programming approaches, using IBM Operational Decision Manager. Decisions and rules for responding to a "probable over-payment" score can be developed using ODM and deployed to the rules engine, where they are accessed by payment processing application programs. This would not require some kind of overhaul of the entire claims processing application. ODM would be used in the development of the new application functionality that would handle claims identified as being at high risk of over-payment.

A workers' compensation use case

When analytics draws on non-traditional data sources, it can enhance the effectiveness of business processes. Consider a workers' compensation system, for example. Suppose an individual files a claim seeking compensation for a work-related back injury. What if, a few days later, the same individual posts on a social media site a picture of himself on a snowboard at a ski resort? (The supposedly injured back would seem to be healthy in that case.)

The worker's compensation system can have a process that pulls information from social media sites using text analytics, with resulting data written to a Hadoop-managed data store. The core application prepares a list of individuals eligible for compensation based on submitted claims, and writes that information to a table in a DB2 for z/OS database.

At that point, a process developed using Spark accesses the compensation-eligible information in the DB2 for z/OS database and matches that with information from social media sites that resides in the Hadoop-managed data store. This augmented data—which includes information about the previously mentioned, and suspiciously timed, snowboarding post—is sent to the claims payment engine, where a more informed decision can now be made on whether to pay a claim or initiate a fraud investigation.

Put real-time analytics technology to work

An enterprise may well have data of great importance residing in z Systems databases, which is currently only delivering a portion of its potential value to the organization. Real-time analytics technology that exists today can boost the return on critical enterprise data assets by infusing data analysis into core application processes, raising a company's capabilities in terms of agility, efficiency and effectiveness. The enterprise data is there. Put it to work in real-time—on the platform where it lives—to drive powerful analytics insight across the organization.

Why IBM?

To learn more about real-time analytics, please contact your IBM representative or IBM Business Partner, or visit: **ibm.com**/analytics/us/en/enterprise

About the authors

Robert Catterall
IBM Senior Consulting IT Specialist, z Systems Data and Analytics

IBM Corporation
IBM Analytics
Route 100
Somers, NY 10589

Produced in the United States of America
October 2016

IMW14879-USEN-00