

# Applying Analytics to IMS Data Helps Achieve Competitive Advantage

Kyle Charlet

Deepak Kohli



 Analytics

z Systems





## The challenge to performing analytics on enterprise data

### Highlights

Business intelligence and analytics helps organization use insights that are drawn from their operational data, which gives them a competitive edge. Traditional copying of data from the source platform to data warehouses is costly and can be a security exposure. Copying data ages data out of its productive use and impedes the delivery of business insight. Data warehouses traditionally do not lend themselves to data intensive analytics. The solution is to bring analytics to the data while retaining the quality of service required by organizations. This approach can have many benefits, such as:

- ▶ Reducing time to market thanks to timely business insights
- ▶ Improved return on investment and performance for analytic workloads
- ▶ Reduced cost of applying analytics through IT infrastructure simplification

The paper *Analytics: The new path to value* written by the IBM® Institute of Business Value (in partnership with MIT Sloan Management Review) states: *At organizations in every industry, in every part of the world, senior leaders wonder whether they are getting full value from the massive amounts of information they already have within their organizations*<sup>1</sup>. A little further, this same paper states *Knowing what happened and why it happened are no longer adequate. Organizations need to know what is happening now, what is likely to happen next, and what actions should be taken to get the optimal results. Obviously, business intelligence, big data, and analytic technologies have been identified as the mechanisms to draw insights from all this data*<sup>2</sup>.

Now consider that IBM z™ Systems mainframes are home to approximately 80% of the world's data, a large part of which is IBM Information Management System (IMS™) transactional data. Why IBM z Systems™? z Systems platforms offer high availability, scalability, security, and performance for online transaction processing (OLTP) workloads. And why IMS? IMS is one of the lowest cost transaction and hierarchical database management systems for mission-critical OLTP workloads available today.

By combining these two distinct yet easily intertwined thoughts, you could ask, "How do I make this highly valuable IMS data available to analytics to perform this much needed analysis and derive new insights? There is always a caveat, so how do I ensure the core capabilities and values of IMS and z Systems are retained?"

Most enterprises began their use of business intelligence (BI) and analytics by first copying data from disparate sources into data warehouses and then typically BI and analytics were used against these copies of the data. Issues emerged from moving copies of the data away from the master source and it was costly to move data across platforms. Plus, this approach removes data from its productive use, which causes the data to age quickly. This situation impeded the delivery of truly valuable and timely business insights. Consider also that where this transactional data originates plays an increasingly important role in analytics. A percentage of the data is not being analyzed at all because of performance issues and the cost challenges of copying data.

<sup>1</sup> *Analytics: The new path to value - How the smartest organizations are embedding analytics to transform insights into action*, found at: <http://public.dhe.ibm.com/common/ssi/ecm/gb/en/gbe03371usen/GBE03371USEN.PDF>

<sup>2</sup> Ibid.

Bringing analytics directly to the data rather than moving data to analytics is the correct answer. It preserves quality of service to ensure timely, accurate, and secure insight and greatly reduces the cost of copying data and storing data. Organizations are realizing that the insights that analytics provide must be delivered efficiently across the organization, with high quality and proper governance.

IBM DB2® Analytics Accelerator for z/OS® is the solution that brings analytics to the data. It is a high-performance appliance that integrates IBM PureData® technology (formerly IBM Netezza®) and z Systems technologies. This solution delivers extremely fast results for complex and data-intensive DB2 queries on various workloads.

This IBM Redbooks® Point-of-View publication describes how DB2 Analytics Accelerator for z/OS can be used to derive business insights and analytic results from IMS data. It also describes the business value and reasoning behind the use of DB2 Analytics Accelerator for z/OS. Use cases in this publication provide examples of the value add this solution can have in various business settings. A high-level description of the steps that are used to make the data available for analysis is provided. Finally, this publication challenges you to look to the future and consider how analytics on IMS data can enrich your business.

## Delivering comprehensive real-time analytics on z Systems

In today's economic climate, analytics' results support fast action, critical thinking, and real-time decision making. The latency of copying data to a data warehouse and running long running analytic queries can make a business decision that is based on this analysis less profitable and timely. How can analytics be performed with current IMS operational data? The answer is two-fold:

- ▶ Bring analytics directly to the data, rather than moving data to the analytics. This approach reduces the cost and security exposure of moving the data.
- ▶ Instead of shoehorning IMS data into a data warehouse for analytic purposes, use an appliance that is built to perform analytics directly on IMS data.

IBM DB2 Analytics Accelerator for z/OS is a built for a purpose, high-performance appliance that integrates IBM PureData technology and z Systems technologies. This combination delivers extremely fast results for complex and data intensive queries on data warehousing, business intelligence, and analytic workloads. This appliance accelerates complex queries up to 2000 times faster. Typical DB2 Analytics Accelerator eligible queries might run an order of magnitude (or more) faster than they would if they used z/OS resources. Queries that are running on the DB2 Analytics Accelerator are not using z/OS resources, which could be a cost savings while allowing other z/OS workloads to perform more efficiently. DB2 Analytics Accelerator for z/OS supports DB2 data and equally important, this solution can be used for all z Systems data sources, including IMS, DB2, and VSAM data.

Other benefits of the DB2 Analytics Accelerator include:

- ▶ No application rewrites or new database designs are required.
- ▶ It uses simplified database design, which results in no need for indexes.
- ▶ The effort to tune analytic queries is eliminated.
- ▶ The DB2 for z/OS data security attributes are inherited.
- ▶ The High Performance Storage Systems (HPSS) feature improves access to and lowers the cost of storing, managing, and processing historical data.
- ▶ Latency is minimized.
- ▶ Fast deployment and quick time-to-value are inherent to the solution.

With IMS operational data made available to the DB2 Analytics Accelerator for z/OS, comprehensive analytics can be delivered to decision makers with industry-leading security and availability. This solution takes the risk out of operational business analytics. The following use cases show how the use of DB2 Analytics Accelerator for z/OS provides organizations with real-time insights.

### **Use case: Being responsive to clients**

One measure of success for any organization is how responsive it is to its clients, whether they are business partners, external clients, or internal organizations. Frequently, clients request customized reports, sourced with data from IMS, DB2 for z/OS, or both. If the necessary information for the report cannot be efficiently gathered, the request cannot be satisfied in a timely manner, which puts the relationship with the client at risk. Each report could require a unique set of data, which causes the need to develop, test, and deploy applications in pursuit of the required data from multiple subsystems. This approach would be costly and time-consuming. In addition, the efficiency of the reporting (running the application) is diminished because of the amount of data that is being processed. In some cases, resource limits could be reached, which can cause the queries to fail. Accessing IMS and DB2 operational data by using IBM DB2 Analytics Accelerator for z/OS creates a single, high-speed approach to analysis with no other application development required to create these reports dynamically, which impresses clients and minimizes impacts to the organization.

### **Use case: Personalized customer care**

Being responsive is a cornerstone of customer retention. The ability to personalize communication and offer customized services (another important cornerstone) helps clients feel like respected individuals. Analytics plays a key role in identifying customized services and finding new opportunities. Data warehouses and repositories depend on data from operational systems, such as IMS (which store transactional detail). Keeping the data volume in these storage facilities under control means much of the detail data must be sacrificed. This approach limits the analytics to working only with the summary data points instead of the full wealth of transactional data that is in the operational data systems.

The use of DB2 Analytics Accelerator for z/OS to perform analytics on the current IMS z Systems-based transactional detail can provide insights that enable you to enrich your customer retention offerings. In addition, organizations that choose to archive the deep transactional detail often maintain that data in managed files, which makes searching and filtering the data costly, complex, and time-consuming. With DB2 Analytics Accelerator for z/OS, this sensitive, historic, and valuable data can be retrieved and analyzed immediately and efficiently.

### **Use case: Data integrity**

IBM z Systems clients frequently have data assets that are housed in IMS and DB2 databases. A correlation often exists between the versions of data in these two data management systems (DBMS) to ensure data integrity. This correlation (a form of referential integrity) exists between the IMS and DB2 database and is often managed at the application layer. In this situation, more applications (usually batch processes) must be run at periodic intervals to ensure that the data is synchronized between the two subsystems.

With the DB2 Analytics Accelerator for z/OS, IMS and DB2 data can be in the same data store, thus minimizing the effort that is associated with writing and maintaining the applications whose sole intent is to ensure data consistency.

## The implementation behind real-time analytics on z Systems

There are three basic implementation steps to make IMS data available to DB2 Analytics Accelerator for z/OS to perform real-time analytics:

- ▶ Extract and optionally transform data
- ▶ Define the IMS metadata
- ▶ Directly load the data into DB2 Analytics Accelerator for z/OS

The key differentiator is that the IMS data does not have to be stored in or routed through IBM DB2 for z/OS. Only the metadata is defined to DB2 for z/OS. Let's look at these steps a little more closely:

### 1. Extract and transform IMS data.

Extract, transform, and load (ETL) tools, such as IBM InfoSphere® DataStage®, can be used to perform the extract step. Homegrown applications can be used to extract data. Some IMS data (for example, IMS data that is stored as packed decimal) can require transformation.

### 2. Define IMS database metadata to DB2 for z/OS and DB2 Analytics Accelerator.

The IBM IMS Explorer for Development, which is included with IMS, can be used to generate the data definition language (DDL) that can be run on DB2 for z/OS to define the IMS metadata. First, the IMS segments are defined in DB2 as tables. These tables can then be defined to DB2 Analytics Accelerator by using IBM DB2 Analytics Accelerator Studio ACCEL\_ADD\_TABLES stored procedure.

### 3. Load the IMS data into DB2 Analytics Accelerator.

Load the IMS data with the IBM DB2 Analytics Accelerator Loader for z/OS tool. You can also optionally load the IMS data into DB2 for z/OS by using the same DB2 Analytics Accelerator Loader for z/OS tool. However, if you intend to store the IMS data only in DB2 Analytics Accelerator, you can use the DB2 for z/OS Accelerator Only Tables (AOT) capability, which tells DB2 for z/OS that the table does not have data.

After these steps are completed, queries that are submitted to DB2 (not to IMS) can be offloaded to DB2 Analytics Accelerator for z/OS, as shown in Figure 1.



Figure 1 DB2 Analytics Accelerator for z/OS interacting with IMS and DB2 data

After the steps are completed, you are now in great shape to begin harnessing the true value of your mainframe data assets by using the DB2 Analytics Accelerator.

## What's next: How IBM can help

IBM is pleased to offer its newly minted Deployment Project Office (DPO) program to help get you started on your analytics journey with IMS. Put simply, the mission of the DPO is to directly engage clients with the IMS lab experts to work jointly towards the goal of production deployment. Although there is no cost for the DPO offering, DPO is exclusively for engagements where production deployment is the goal. For more information, contact us at this email address: [ibmims@us.ibm.com](mailto:ibmims@us.ibm.com).

In addition, you can contact your IMS advocates who can also help with general questions and proof of concepts that are starting or underway.

## Resources for more information

For more information about the concepts that are highlighted in the paper, see the following resources:

- ▶ IBM DB2 Analytics Accelerator for z/OS product page:  
<http://www.ibm.com/software/products/en/db2analacceforzos>
- ▶ *Accelerate business insights with IMS transactional data*:  
[http://ibm.biz/accelerate\\_insights\\_ims\\_transactional\\_data](http://ibm.biz/accelerate_insights_ims_transactional_data)
- ▶ IBM DeveloperWorks documents on Implementing DB2 Analytics Accelerator on IMS data:  
[https://ibm.biz/ims\\_idaa\\_technical\\_implementation](https://ibm.biz/ims_idaa_technical_implementation)
- ▶ *Using the Power of DB2 Analytics Accelerator on IMS Data* YouTube video:  
[https://ibm.biz/demo\\_ims\\_idaa](https://ibm.biz/demo_ims_idaa)





# Notices

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

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

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM 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 supply in any way it believes appropriate without incurring any obligation to you.

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


## 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.

# Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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:

DataStage®	IMS™	z Systems™
DB2®	InfoSphere®	z/OS®
IBM®	PureData®	zEnterprise®
IBM PureData®	Redbooks®	
IBM z™	Redbooks (logo)  ®	

The following terms are trademarks of other companies:

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

Other company, product, or service names may be trademarks or service marks of others.





REDP-5295-00

ISBN 0738454583

Printed in U.S.A.

Get connected

