Making data simple and accessible

The role of technology in delivering analytic results

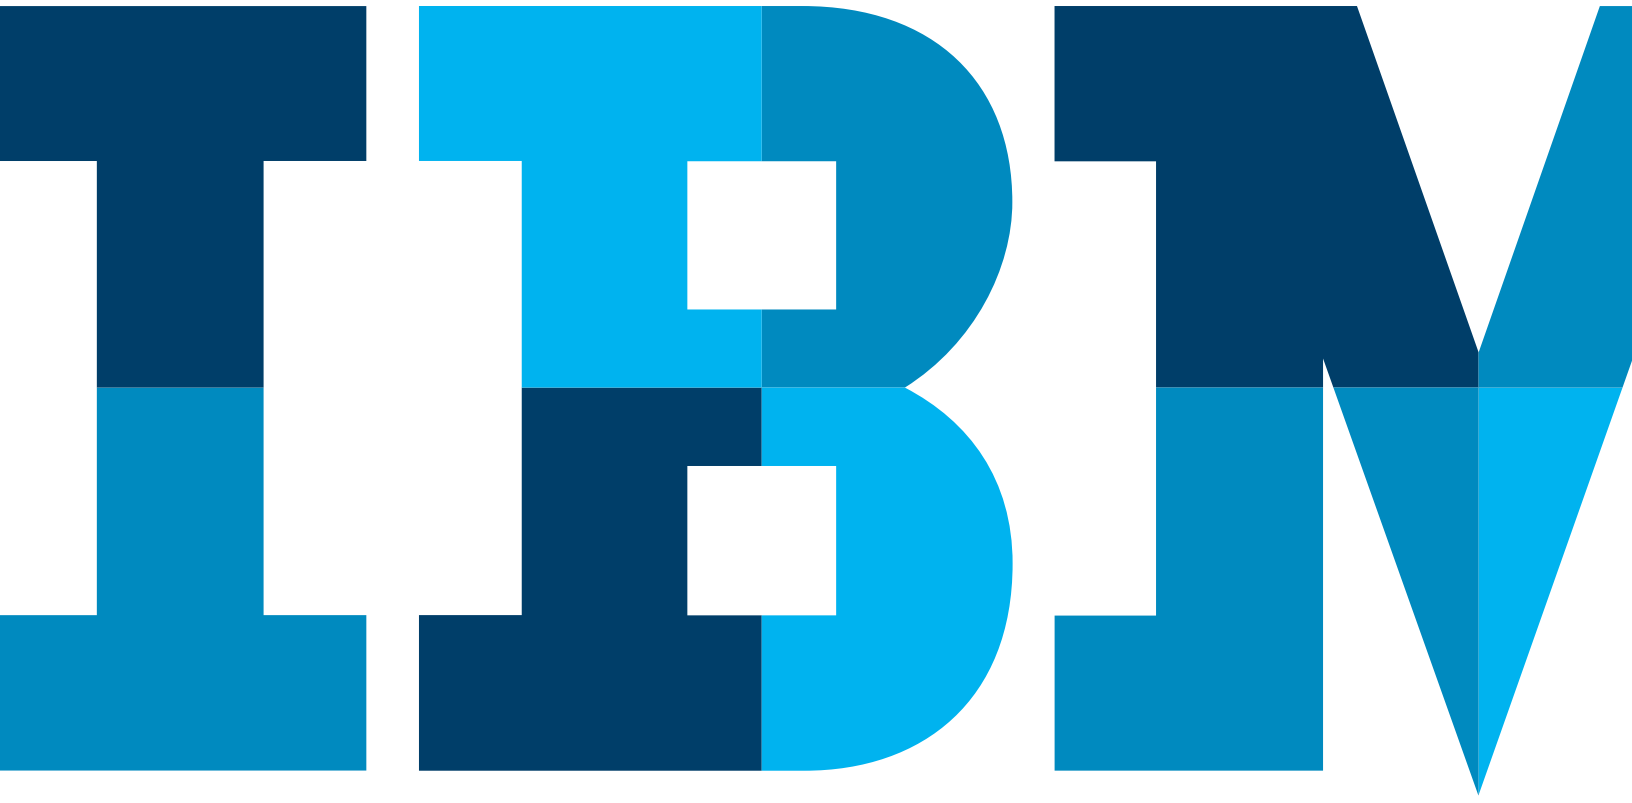

IBM

Contents

Overview

The vast amount and variety of data available to today's enterprises provides significant business opportunities. Increasingly, organizations have access to multiple sources of structured and unstructured data that can lead to keen insights on aspects of the business, which can provide a competitive advantage. Moreover, the capability to derive the necessary insight from that data to rapidly make critical business decisions is key—that's when analytics come into play. An essential component of analytics processing is query workload execution. The need to get information increasingly faster than was possible before and to enhance the performance of query execution for 24x7 demands creates new challenges.

A technology revolution on z Systems

The introduction of the IBM® DB2® Analytics Accelerator, in conjunction with IBM z Systems™, has brought a revolutionary change—resulting in a hybrid computing environment on z Systems.

Figure 1: DB2 for IBM z/OS® and DB2 Analytics Accelerator, a transaction and analytic processing system

This hybrid environment supports transaction processing and analytic workloads concurrently, efficiently and cost-effectively, while delivering industry-leading performance for mixed workloads. The integration is transparent to users and application changes are not required to accelerate the query.

Understanding user roles and data

More organizations are becoming data driven, enabling widespread use of analytics across the enterprise, and collaboration and interaction with data to gain greater insights.

The information and analytic lifecycle engages with different data roles across its path. Historically, the primary focus was on the IT organization who served as the ultimate keepers of the data, providing access to the data. This was more of a bottom-up design. Today we have roles, such as chief data officer (CDO), data engineer, data scientists and business analysts. These positions help push the design and focus of the analytic architecture from the top down through the perspective of the data consumers. This approach allows the data consumers to be more involved in the design, collaborate across the organization and innovate faster.

This white paper focuses on several customer case studies, illustrating how IT makes the necessary data available to specific users for determining analytic results. The addition of the DB2 Analytics Accelerator Loader for z/OS plays a key role for the IT department in achieving the goals specified by others in the organization. It solves many of the challenges in dealing with disparate data types used for analytics.

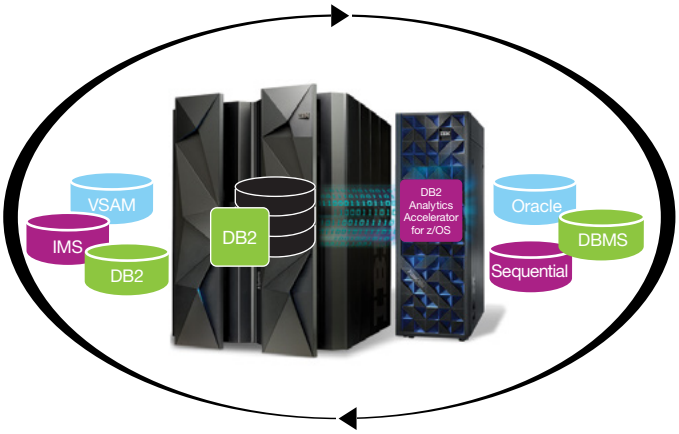

Figure 2: Transforming data for use in analytics

Companies capture a wealth of operational data on z/OS. Bringing analytics to that data can deliver powerful insights to drive key business decisions.

One of the primary challenges companies face is how to combine multiple data sources to deliver the right data to the right users at the right time with the least possible cost or risk. Many companies perform time-consuming and costly extract, transform and load (ETL) operations to move data to other platforms to perform analytics. However, there are potential security exposures and risk with each data movement operation.

The Accelerator Loader solves the IT department's challenges of using different data types for analytics and reduces data latency for the analytics applications and users. The Accelerator Loader is crucial to delivering the analytic results to various lines of business (LOB), and also provides necessary transparency to users and significant cost savings to the organization.

Use Cases

The following section provides several examples of how the combined use of the DB2 Analytics Accelerator and the Accelerator Loader was used to deliver analytic results to users.

Avoid data movement

A large US financial institution was moving Visual Storage Access Method (VSAM) data off platform to perform analytics. Before implementing the IBM Accelerator, it was doing nightly ETL operations to move data off of z Systems. The company estimated that over the past three years, prior to using the Accelerator, it was spending millions of dollars to run its analytics on distributed systems. By leveraging the IBM Accelerator and the Accelerator Loader, the company is able to keep the data on z Systems and accomplish its analytics goals with significantly less data latency and costs [1].

The Accelerator Loader is able to automate the entire process of loading data for use in the Accelerator, including:

- Creation of the DB2 table
- Addition of the table to the Accelerator
- Extraction of the specified data source
- In-memory data conversion to the necessary DB2 format
- Loading data to the Accelerator
- Enablement of the table for acceleration

The savings and benefits are increased by the use of the Accelerator Loader with the elimination of ETL operations and specialized programing costs.

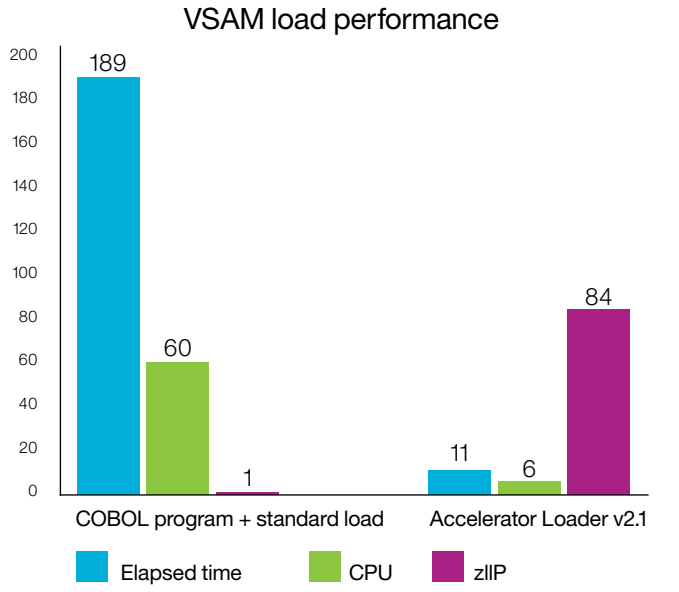

Figure 3: Accelerator Loader example loading 200 GB VSAM data with six z Systems Integrated Information Processors (zIIPs) using a simple COBOL program with the standard load operation to extract the data

Additional benefits gained with the Accelerator Loader, illustrated in Figure 3, include:

- *Significant CPU savings that lowers* million service units *(MSU) pricing*. In this example, the financial institution used approximately 10 percent of the original CPU.
- *Significant reduction in ETL batch window time*. In this example, it went from greater than three hours down to 11 minutes. The company is able to perform analytics closer to real time. The ETL savings does not include the time and effort to develop, test and maintain the ETL COBOL program.
- *Greater utilization of z Systems*. In this example, the company operated an IBM z13™. More than 80 percent of the VSAM Loader process is eligible for the zIIP. The Accelerator Loader also exploits z13 multithreading for an additional 50 percent zIIP offload.

Integrate multiple data types

A large worldwide financial institution originally invested in the DB2 Analytics Accelerator to reduce queries from hours to seconds for a large digital campaign it was running. The inclusion of the DB2 Analytics Accelerator Loader allowed it to re-architect current business analytics applications to enable loading of non-DB2 data from other platforms. The company stated that the DB2 Analytics Accelerator Loader opened up more data types, reduced data latency and provided additional storage savings [2].

This financial institution is now able to reach out to more of its foreign affiliates and take advantage of greater analytics capabilities with multiple data centers. The users of the worldwide marketing campaign have no idea that they are accessing and joining data from Europe to the US. Since the results from user queries do not have to be stored, they can be reissued at a later time, if needed, to compare the results. The other big advantage to users is that they don't have to worry about security or data privacy issues. Many users attempt to run ad hoc queries without thinking about whether they have the necessary authorization or not. Existing access control in the backend system is used to prevent unauthorized access to the data. The financial institution's senior IT specialist claimed that with the DB2 Analytics Accelerator and Accelerator Loader, "Data privacy is not an issue with its real-time data and analytics processing."

Additionally, the company developed an iPhone app that allows clients to check their 401K balances at any time and obtain updated, instant information. The company views this application usage as an overall revenue enhancement to its business. It also is aimed at a generation who expects information available at their convenience without having to wait or deal with a voice-recorded phone application.

The company stated that the combination of Accelerator and Accelerator Loader is truly real-time analytics. To date, this solution has yielded it a 70 to 80 percent reduction of elapsed time for its critical business applications and it expects even greater results in the near future.

Run analytics on operational data

There are many types of data produced by z/OS and other installed software that is stored on z/OS. IBM System Management Facility (SMF) is IBM's data collection facility for accounting and performance data on z/OS. Different SMF records are created for all functions within the z/OS system and its related subsystems. There is also IBM Resource Measurement Facility (RMF) data, various performance metrics collected from software or customer programs, and log information produced by DB2, IBM CICS® and IBM WebSphere®. All of these data types may be categorized as operational data.

The chief data officer and IT executive team of a worldwide insurance company wanted to use analytics on operational data to provide insight into its system usage for accurately pinpointing system needs on a timely basis. However, the company could only store 30 days of SMF data on z/OS. There were situations where the system reached full capacity levels and the company had to delay new projects. And conversely, purchases made in anticipation of growth did not materialize in the fiscal year.

For the IT department, capacity planning activities involved analyzing years of SMF data. The company was able to first move SMF data from z/OS storage to the Accelerator's high-performance storage saver (HPSS) and store considerably more than 30 days of performance data. Analytics were used to provide timely and accurate capacity planning and performance metrics that could help the company better estimate when and where additional resources were needed. The company realized savings from using HPSS to store the data. It was also able to eliminate existing software that had been used for capacity planning, resulting in greater ROI from its analytics investment[3].

Summary

The DB2 Analytics Accelerator — together with DB2 for z/OS and z Systems — form a self-managing, hybrid workload-optimized database management system. This hybrid solution runs each online transaction processing (OLTP), batch and analytic query in the most efficient way, enabling each execution in its optimal environment for greatest performance and cost efficiency. The addition of the DB2 Analytics Accelerator Loader opens the data aperture for greater analytic insights, while reducing complexity, latency and costs.

Through this hybrid solution, IBM can bring high-volume business transactions, batch reporting and complex analytic queries together and run them concurrently in a mixed-workload environment. Bringing your analytics closer to your transactional data source can help reduce data latency, complexity and cost. It can also help you deliver your business-critical analytics with z Systems qualities of service to provide more timely, accurate and secure data; superior availability, scalability and performance; and rapid deployment and expansion. So you can confidently apply those same qualities of service that you depend on today, from your transactional systems to your analytics. This hybrid approach provides a flexible platform that enables you to start with your top analytic requirements to quickly realize business value. You can position your analytics strategy to grow and evolve as your business and the market demands dictate — without the need to change your existing IT environment.

For more information

To learn more about IBM DB2 Analytics Accelerator and DB2 Analytics Accelerator Loader, visit:
ibm.com/software/data/db2/zos/tools

IBM Corporation Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
July 2016

It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. 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.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NONINFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

1 Based on IBM internal data analysis. Individual client results will vary.

2 Based on IBM internal data analysis. Individual client results will vary.

3 Based on IBM internal data analysis. Individual client results will vary.

Please Recycle

IMW14876-USEN-00