WHITE PAPER

PERFORM REAL-TIME ANALYTICS & MACHINE LEARNING ON IBM Z

IBM Silver Business Partner™

Infosys® Navigate your next

Abstract

In the era of rapid digitization, business expects personalized customer interaction and real-time analytics as an essential experience. 80% of B2B customers and 64% of consumers expect companies to respond and interact with them in real time [*]. We are seeing a dramatic growth in advanced analytics market from the arrival of Big Data and connected devices to the development of AI-driven predictive, cognitive branches.

As 80% of the corporate data is in mainframe, most of the businesses are looking to integrate mainframes into a broader analytical environment. This white paper explores various emerging trends on how advanced analytics and machine learning are emerging for data in big iron and how to take advantage of the options available.

*https://www.ibm.com/it-infrastructure/z/capabilities/real-time-analytics

Introduction

As of 2018, mainframes process around 30 billion business transactions per day, including most major credit card transactions and stock trades, money transfers, manufacturing processes and ERP systems. It is estimated that 80 percent of the world's corporate data resides or originates on mainframes [*]. Ten out of top ten insurers process their high-volume transactions on mainframes and 92 of the top 100 global banks run on mainframes. These impressive statistics for today's mainframe adoption are a direct counterpoint to the misperceptions in some circles that the "mainframe is going away".

For large organizations who have heavy processing on mainframe (for example, banks, insurance providers, healthcare industry and retail giants), moving out of mainframes is still not a viable approach. Their current focus seems to be on modernizing their mainframes to support their digital journey through emerging trends like API enablement on mainframes, embracing Agile/DevOps etc. The volume of core business processing that happens in their mainframe environment is so critical and significant in volume that it becomes extremely challenging to match the mainframe's throughput, reliability, availability, security and serviceability.

According to an IDC Info Brief, sponsored by Broadcom Inc. and IBM, conducted in 2019, the mainframe has evolved from siloed to connected to transformative and with that the mainframe platform is transforming from a revenue-supporting machine into a revenue-generating machine and is increasingly playing a central role in organization's digital transformation (DX) journey. The IDC paper finds that the modernization and integration initiatives for achieving transformative mainframe leads to new business innovations, which in turn are driving revenue growth and improving organizational operational efficiency. As per the IDC study, adopters of a transformative mainframe strategy can achieve average benefits worth more than 6x what they invest to transform their mainframe platforms. By improving business and IT staff productivity and cutting operational costs, they can generate an average of almost $200 million in additional revenue per year. IDC calculates that the adopters of transformative initiatives reduce hardware and licensing costs by 14% on average. IDC finds that at this level of mainframe management and cost efficiencies, the adopters will reduce the cost of operating their mainframe platforms by an average of 19% over five years through their mainframe transformation efforts. IDC study shows those who adopted the transformative mainframe approach has brought down already low levels of unplanned outages by an average of 43%. The study also shows the adopters of transformative mainframe can achieve 64% more code releases and 30% more efficient mainframe management and thereby significantly improves the business agility.

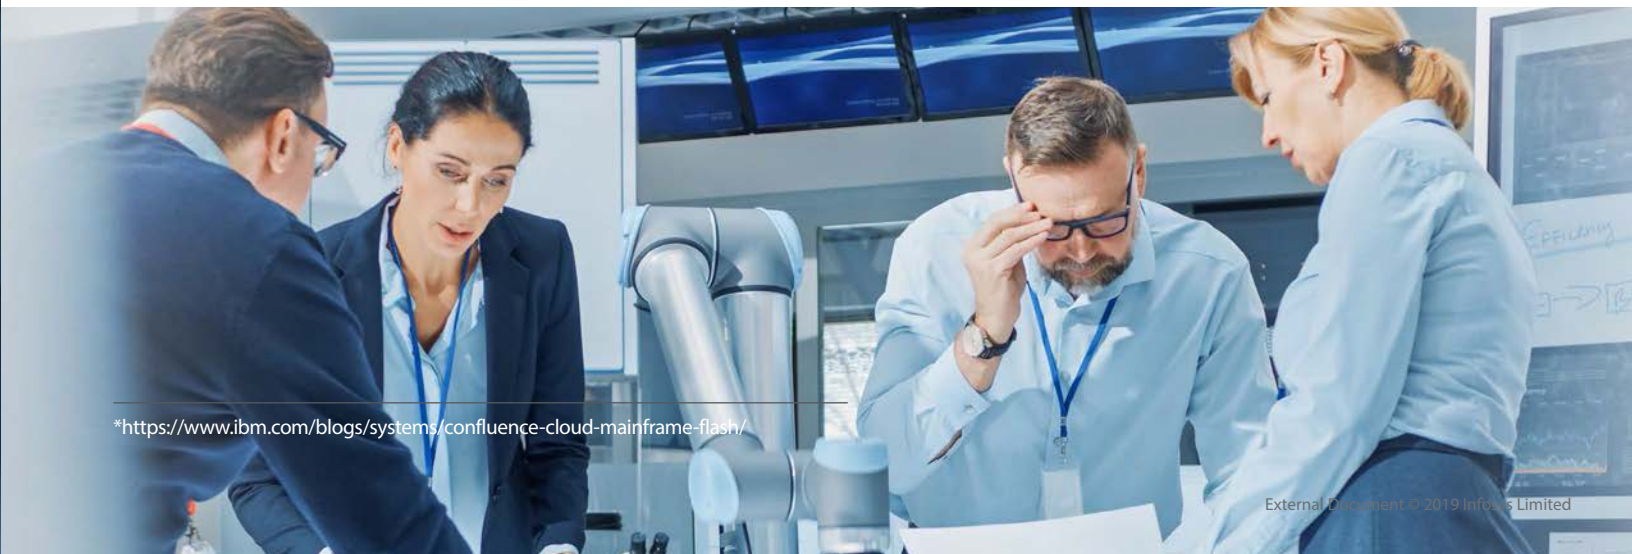

Source: *The Business Value of the Transformative Mainframe for Digital Transformation, IDC white Paper, sponsored by Broadcom and IBM, August 2019* *https://www.ibm.com/account/reg/us-en/signup?formid=urx-4008*

*https://www.ibm.com/blogs/systems/confluence-cloud-mainframe-flash/

Similarly, a BMC Mainframe Research Survey conducted in 2018 also depicted how digital business is driving growth in mainframes according to the technology leaders across the industry. Digital business is driving growth in the mainframe in five key categories.

Mainframe Growth

According to the survey, Digital business is driving growth in the mainframe in five key categories

This activity level has driven 60% of the technology leadership to have a positive outlook for MIPS growth—the highest level in three years.

Breakout – Large Shops

Large shops [>10,000 MIPS] outplaces all respondents with their increases.

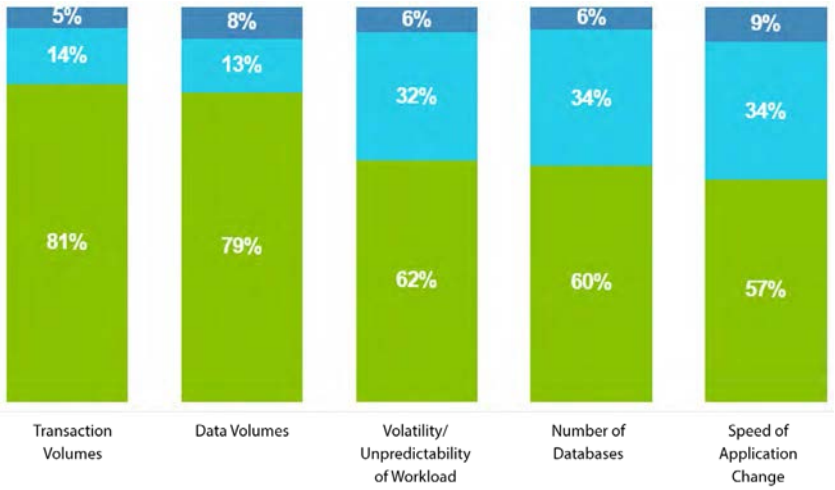

This activity level has driven 70% of large shops to have a positive outlook for MIPS growth.

Survey respondents answered as

Increased **Stayed the same** **Decreased**

Considering the growth rate in data and transactions volume, the field of data analytics, AI and Big Data are also evolving at rapid pace to support the digital growth. The IDC analysis to find the business value of Transformative mainframe also shows the growing interests of enterprises for doing the analytics and AI on mainframe.

* https://www.bmc.com/content/dam/bmc/collateral/bmc/Annual-Mainframe-Survey-Report-2018-v8.pdf

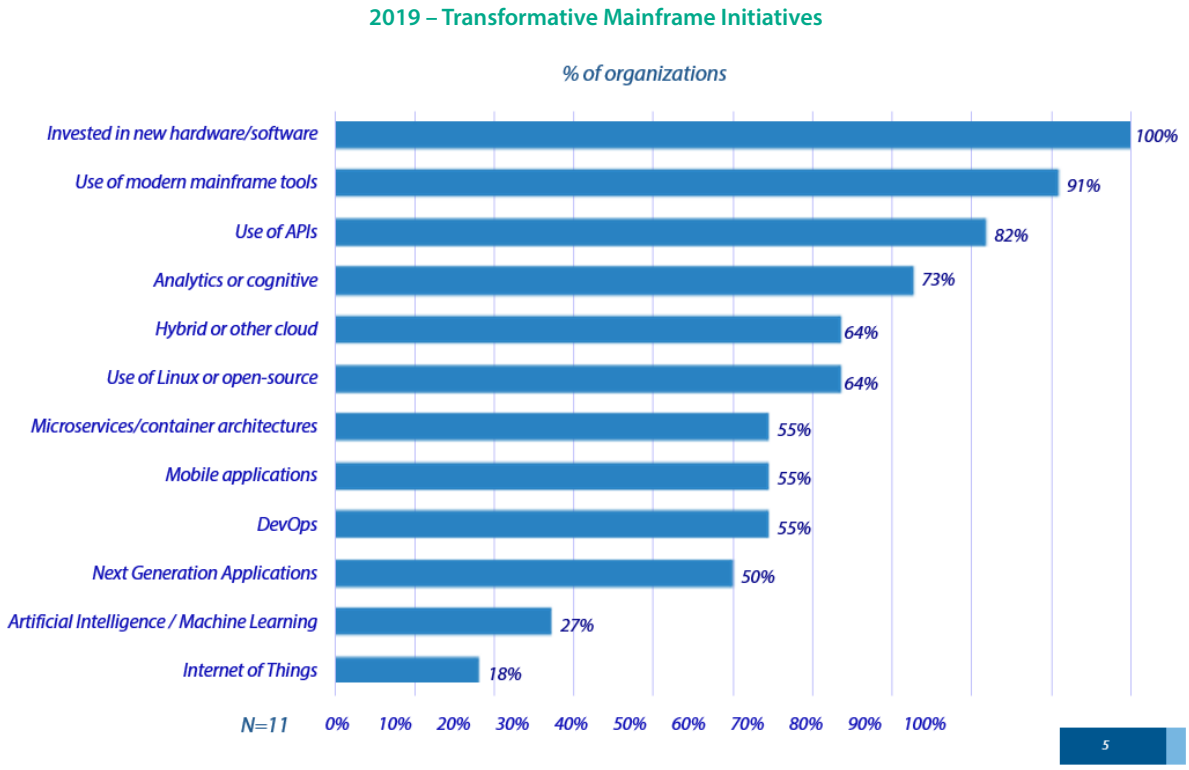

Source: *: The Business Value of the Transformative Mainframe for Digital Transformation, IDC white Paper, sponsored by Broadcom and IBM, August 2019*

- **73% wants Analytics or Cognitive**
- **27% wants AI/Machine learning**

Among the enterprises participated in the survey, 73% of the mainframe customers are interested to do the analytics on mainframe whereas 27% show interest to do AI on mainframe.

Given these trends, one should consider the option to run data analytics directly at the source of data that is many times on mainframe systems, particularly if these analytics can be done ably and cost effectively.

Before we discuss on the options available to implement data analytics and machine learning on mainframe, let us understand how data analytics on transactional data is typically happening now.

Current State and Challenges

Most data centric organizations have a defined process in which they ingest data from various sources, store the ingested data into the data lake, process the data and then build the reporting and visualization layer on top of it. The process involves ETL (Extract, Transform, Load) solutions, where the data is gathered from multiple sources (primarily mainframes), structured, organized and centralized into a single repository such as a data lake.

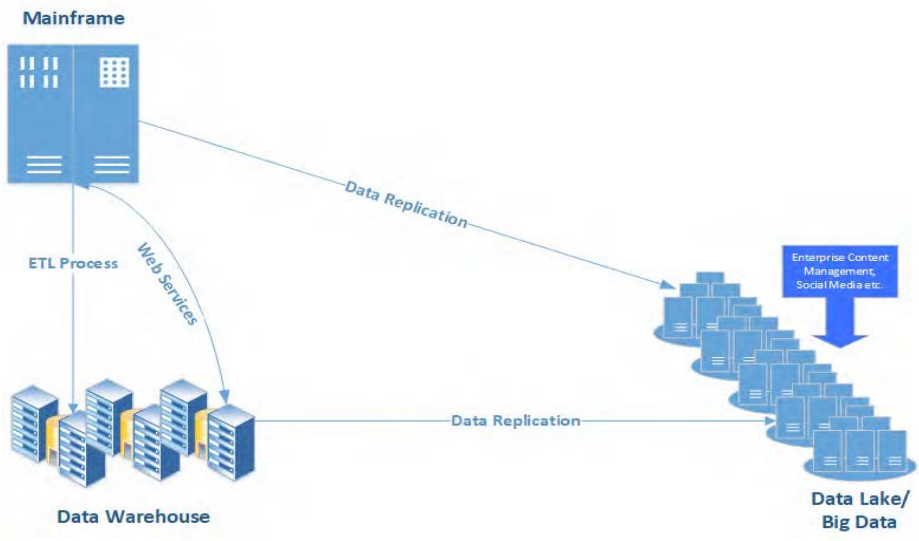

A typical 3-cluster structure of data congregation observed today is

1) *Mainframe acts as a transactional backbone (system of record)*
2) *Data Warehouse for Business Intelligence and*
3) *Data Lake assimilate both structured and unstructured data.*

The challenges we see in the current approach are:

- *Latency caused by any overnight batch and ETL process act as a roadblock for real time requirements of business units (e.g. in ecommerce order fulfillment, insurance products purchase, fraud and risk detection for payments and others etc.). Almost all critical data generally is moved out of mainframe, which introduces data redundancy, data sync/replication issues and time lag.*
- *ETL process runs either batch process or started tasks in mainframes, which adds to the overall mainframe software licensing cost*
- *Any solution to overcome latency increases the overall cost to manage and administer because of over engineering*
- *Complexity, security implications and privacy needs of the data movement and storage is constantly increasing*

The emerging data types like click stream data, social media unstructured data, etc. are processed by systems outside mainframes. To perform a holistic data analysis and implement machine learning on mainframes, it is important that the data from external systems are also accessible from mainframes. Therefore, every enterprise that has significant portion of its system of records (SOR) on mainframes, encounter challenges as to visualize, plan and construct the organization's data analytics platform that support critical decision-making uses cases constructed on applications both on and off mainframes.

Factors Influencing Implementation of Analytics/ML on Mainframe

According to our experience, there are multiple evaluation points to consider before an organization invests on a platform that enables data analytics or machine learning on mainframe. Let us analyze these factors in detail.

1. *What share of real time business decision making happens in mainframe?*

For any organization, growth of business would have made their applications big and diversified over the years. Say an application that was first developed on mainframe (closed and monolithic), would have evolved and became an application adopting a 3-tier architecture (UI Layer, Business Layer and Data Layer). If the application's business layer and the data layer is still on mainframes, it makes a good case that the business decision making is still running on mainframes. The more the business logic and transaction data is on mainframes, the more data gravity. *(*Data gravity is an analogy of the nature of data and its ability to attract additional applications and services. Dave McCrory coined the term data gravity to describe the phenomenon in which the number or quantity and the speed at which services, applications, and even customers are attracted to data increases as the mass of the data also increases.)*. In addition to the considerations about application structure, the value of current insights to the use case should also be considered. If only ETL'd data is used for analysis, then the insights generated can be based on older, stale data. The value of current insight and in some cases real-time insight can provide business benefit. In these cases, analytics at the data source where it originates, such as mainframe, can be a critical capability.

2. *Is the data moved out of mainframes primarily to run data analytics?*

Most large enterprise who use mainframes have built data warehouse solutions where the data from mainframes is also available in structured format for other system consumptions and archival needs. However, if the data is moved outside mainframes primarily to run analytics over the extracted data, it may no longer be required given the recent innovations on the mainframe and offload approach is costlier over time.

3. *Security and Sensitivity of the data*

Mainframe is known for its high security and it is by far the best platform that has been time-tested with respect to avoid data breaches. Considering this salient feature, it is ideal to restrict the movement of sensitive information from the mainframe to less secure environments. Compliance to security Regulations is another important factor to consider while moving secure data out of mainframe. At the same time, it is also important to note that restricting data movement to outside systems is not always possible.

4. *Size/Volume of data to be offloaded from mainframe for the analysis:*

Size of the required data for a given use case should also be considered. If the data size is significant, that may cause both inefficiencies in processing as well as increase in cost. The mainframe environment has capability to run a wide range of analytics natively on the platform and share results through REST based APIs. This kind of hybrid approach can result in agility and flexibility for clients.

5. *Amount of data transformation needed to make it useful to perform data analytics*

IMS hierarchical database and VSAM datasets are still prominent in the mainframe even after RDBMS solutions like DB2 exist. It is mainly because some businesses have built their primary data so complex and vast in IMS that it takes huge effort and high risk to migrate them to RDBMS tables. Therefore, in situations like these, effort and tools/software are required to extract, transform and load the data from mainframes and ingest into the data warehouses. Leveraging data-in-place analytics strategies for mainframe data (IMS, VSAM, DB2 z/OS, PS, etc.) can streamline the use of this valuable data for analysis, resulting in shorter time to realize the business results.

Data Analytics and Machine Learning on Mainframes

Enterprises' mainframes have a wealth of transactional, operational and customer data. Most enterprises currently depend on their data warehouse to run their data analytics and model any machine learning solutions. Depending on the data latency, some of the insights are perishable, and it may lose its relevance if the data synchronization frequency between mainframe and data warehouse is wider. We can make more relevant business decisions if you could capture these insights in real time.

It makes more sense to co-locate analytics securely with data based on volume, rate of change and value. In practical terms, moving data farther and more frequently affects workload performance. With Analytics on z/OS, batch workloads like reporting, consolidation can bring near real-time insights for traditional batch processes.

Let us have a look at the key analytics aiding products at the IBM Z analytics portfolio.

IBM Open Data Analytics for z/OS [IzODA]

IBM Open Data Analytics for z/OS exploits in-memory computing while leaving continuously updated data securely in place to improve analytics accuracy and accelerate predictive behavior models. This means that we can run any kind of analytic capabilities, including machine learning, using the most current data. According to a recent IBM commissioned Forrester survey, it is not a surprise that 91% of data scientist prefer real time data for modeling.

IBM Open Data Analytics for z/OS primarily consists of four major components:

1. Apache Spark: A high-performance, general execution engine for large-scale data processing. One of its key features is the capability to perform in-memory computing. Unlike traditional large data processing technologies, Spark allows caching of intermediate results in memory rather than writing them to disk, thereby dramatically improving the performance of iterative processing. In addition to execution engine (Spark core), Spark also provides libraries for analytics related functions like Streaming (Spark Streaming), Machine Learning (MLib), SQL processing (Spark SQL) and Graphics processing (GraphX).

Spark Core	Spark Streaming	Stream ProcessingNear real-time data processing & analytics
	MLlib (Machine Learning)	Machine LearningIncredibly fast, easy to deploy algorithms
	Spark SQL	Unified Data AccessFast, familiar query language for all data
	GraphX (graph)	Graph AnalyticsFast and integrated graph computation

2. Optimized Data Layer (ODL): provides integration facilities for both IBM Z data sources and other off-platform data sources. The Data Service provides your Apache Spark, python, pyspark application with optimized, virtualized, and parallelized access to a wide variety of data.
3. Anaconda: Anaconda includes Python and Anaconda Python packages for data science, which provide data scientists with a comprehensive solution for integrating computations to the data.
4. Apache Livy: It is a REST service used in conjunction with Spark that enables users to submit Spark jobs without having the Spark client installed. This enables developers to harness the data analytics power that Spark is capable of providing from within a web or mobile application.

The key benefits of IBM Open Data Analytics for z/OS are the following:

- *Federated analytics across multiple data environments and better integration with enterprise business applications*
- *Direct, real-time access w/o any latency to business-critical mainframe data and other loaded data thereby providing real business benefits that are not attainable otherwise as well as savings in Offload-MIPS, storage, admin efforts*
- *Extremely high zIIP utilization for affordability and low-cost solution*

Clients have several choices to configure IzODA in their hybrid environments. Two of these options are as follows.

1. Install IzODA on mainframes and any data access to smaller data marts or larger data warehouses can be achieved through in-memory processing via Optimized Data Layer

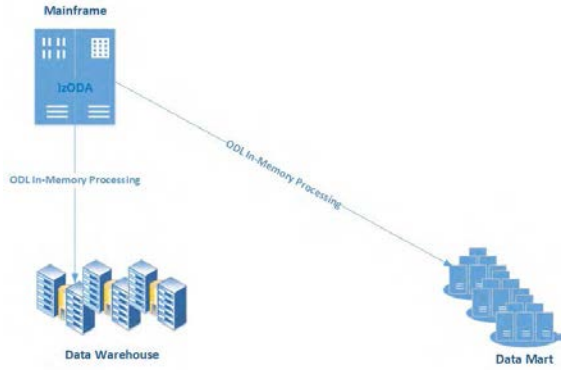

2. Install IzODA on mainframes and keep separate spark cluster for distributed platforms. Do federated analytics by submit the spark job for distributed platform from mainframe through Apache livy server and consolidate the results on mainframe.

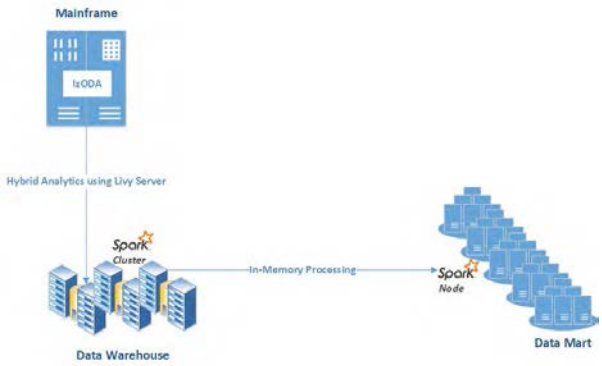

Federated Analytics: IBM Z and Distributed Platform Interaction

The high-level architecture of IzODA for federated analytics between IBM Z and Distributed platforms is depicted in the below picture.

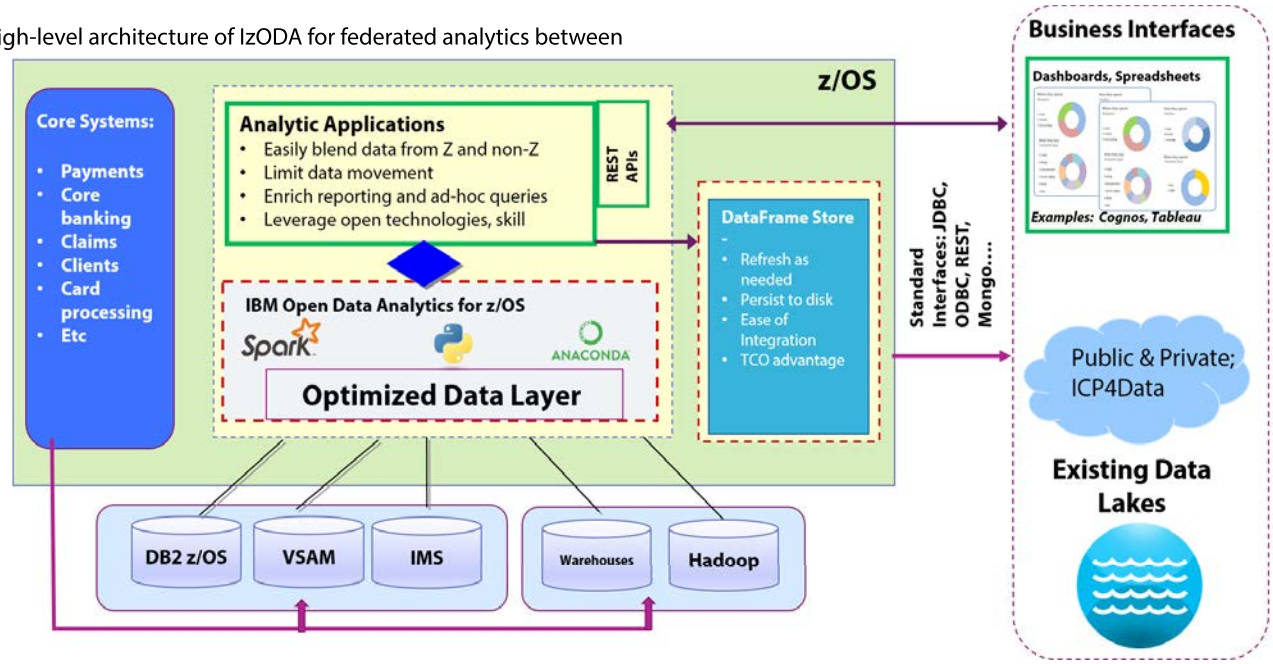

With IzODA, federated analytics is very much possible with mainframe and distributed platforms. No need to move the data between platforms as the data can be analyzed at data source itself. The data from IBM Z and non-Z platforms can be accessed transparently through optimized data layer of IzODA. The analytics results from mainframe can be stored in Data Frame Store of IzODA which can be accessed from distributed platforms through REST, JDBC, ODBC etc. interfaces. As the same open source technologies are used for analytics on all platforms, skill is not an issue to work with analytics on mainframe. The existing reporting tools and dashboard can be leveraged to display the mainframe analytics results by accessing through standard APIs.

Watson Machine Learning for z/OS [WMLz]

Watson Machine Learning for z/OS is a product tailored for machine learning applications where clients want to have an end to end framework for development, training and deployment of machine learning function, governance, monitoring, etc. It includes the analytics runtime (IBM Open Data Analytics for z/OS) as well.

WMLz followed a two-tiered architecture. The analytics runtime runs on z/OS whereas the IDE runs on Linux x86 or Linux on Z.

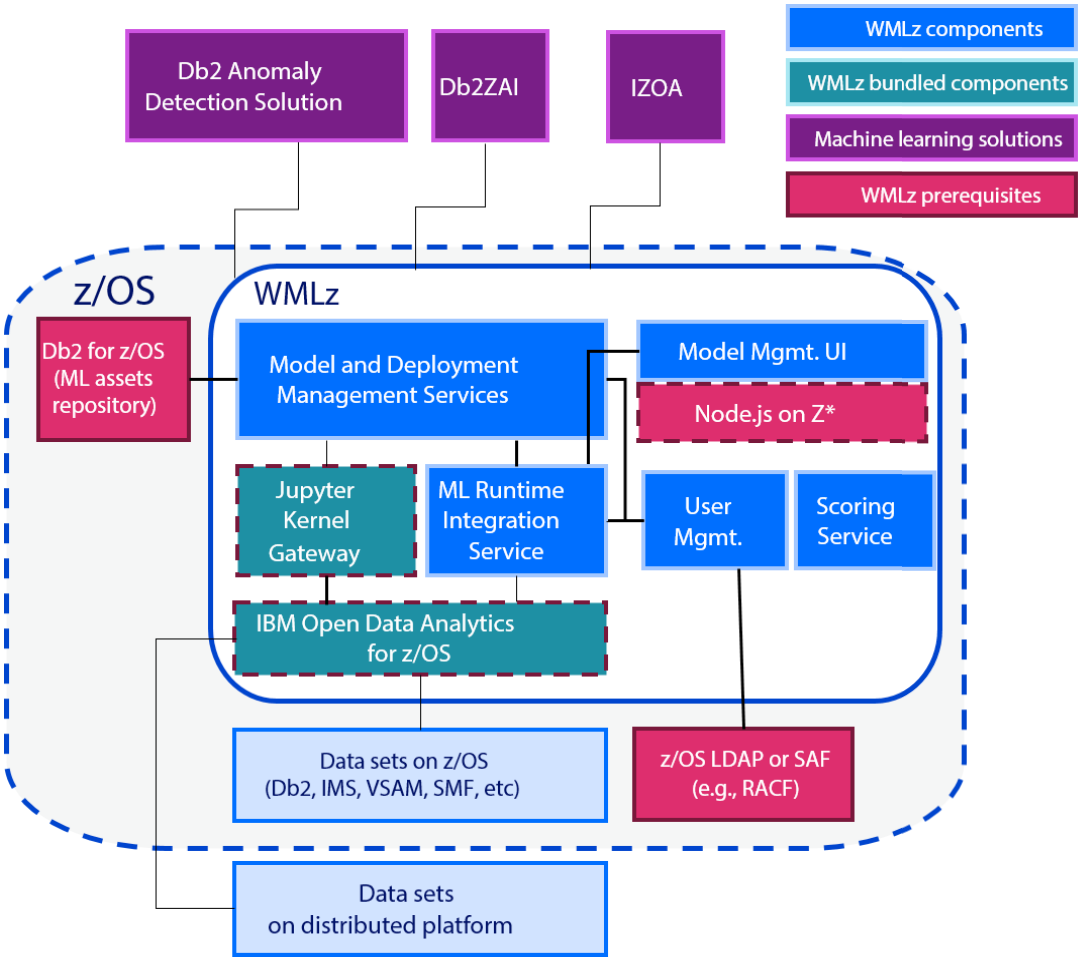

The key features of WMLz are the following:

- **Flexible model development:** Flexibility to build, train and evaluate models using their IDE of choice or the extensive model building features based on enterprise-grade open source software in Watson Machine Learning for z/OS.
- **Improved productivity:** Optimize data scientist productivity through extensive Watson Machine Learning for z/OS model building features/modes including notebooks, visual builders, wizards and enhanced intelligence applied to data scientist activities. Automatically normalize, handle missing values and generate data features to make even novice data scientists into experts.
- **Enterprise-ready AI model deployment:** Operationalize predictive models within transactional applications, without significant overhead, enabling real-time insight at the point of interaction. Offers several scoring approaches including RESTful APIs and Java and CICS integration, optimized for the highest security and performance level on IBM Z.
- **Enhanced model accuracy:** Enable data scientists and engineers to schedule periodic re-evaluations of new data to monitor model accuracy over time and be alerted when performance deteriorates. Automatically refresh models to maintain model accuracy with confidence.
- **Production-ready machine learning:** Deliver essential model versioning, auditing and monitoring as well as high availability, high performance, low latency and machine learning model automation.
- **Quick-start solution templates:** Offer essential foundational templates for common business requirements to bootstrap your machine learning efforts. Solution templates demonstrate how machine learning can run alongside your application infrastructure to add value to key business areas including fraud detection, loan approval and IT operational analytics (ITOA).

Industry Use Cases

Use Case #1: Applying predictive analytics to reduce the number of order change requests in automotive supply chain

Problem Statement:

In the automobile industry, more than 95% of all the vehicles sold are "Built to Order". Although there are lot of inputs taken into consideration during forecasting and ordering (like historical sales data, product mix, fast moving colors, dealer preferences, accessory specifications, etc.), there is significant percentage of orders, that undergo order change process (configuration change). e.g. In US, some warm states such as California or Florida usually sell cars with no cold weather kits. However, other few cold states such as Michigan or Illinois cannot sell cars without cold weather kits. Likewise, dealers in few regions decided that an antilock brake system (ABS) would be an option most customers would not be willing to pay for; thus, vehicles with these systems would not be stocked.

After the confirmed orders are fed into the system, the allocation system allocates vehicles to the dealers. Now the dealers check their allocated vehicles and if they require changes to the vehicle configuration, they can submit change requests within a specific time before the vehicle configuration is locked for production. The dealership personnel, HQ and Regional Order Change Admin submit change request to the vehicles to match a configuration that generally sell fast in the dealer area.

The change requests generally involve three major parameters namely color, model and Factory Installed Options (FIO). Below diagram represents a sample order change data categorized by type for a US automotive major.

Change Request By Type

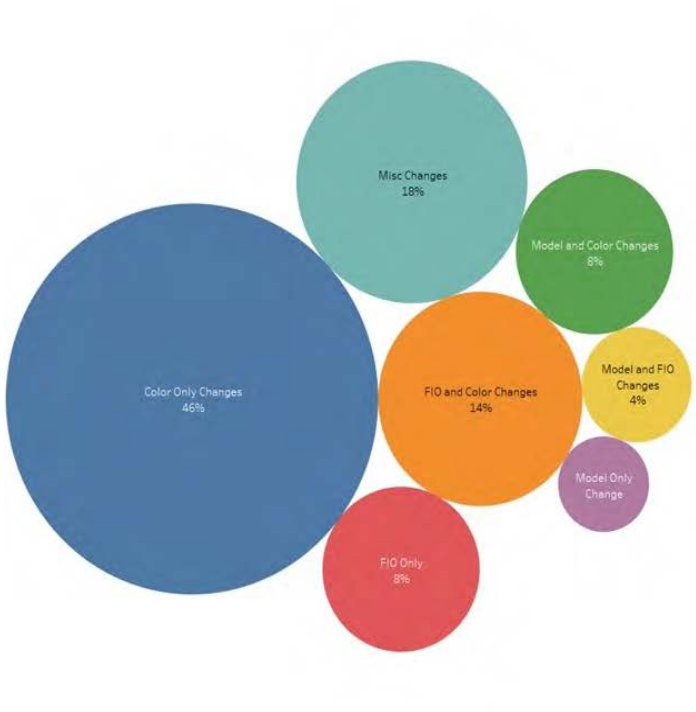

Change Request Acceptance By Type

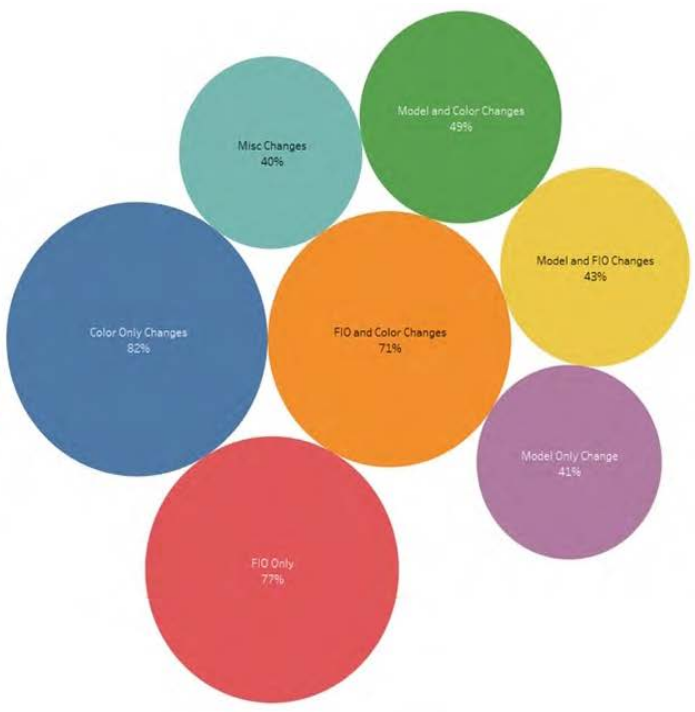

Challenges:

- The OEM dealer requests mass changes to confirmed units prior to build immediately after month end allocation (first 3 days at the start of a month). These change requests are predominantly repetitive in nature
- Due to huge volume of change requests, sometimes the ETL jobs to offload the data from mainframe to data warehouse run for longer time and the requests fail to reach the processing system within plant daily cut off SLA time and in-turn increasing the delay in acceptance response by 1 day

Proposed Solution:

The core processing and data for the order change processing is in mainframes. In the current state, all the core mainframe data are ETLed out to data warehouse where the analytics is performed. Enabling mainframes to support analytics and ML solutions with the help of products like IBM Open Data Analytics for z/OS and Watson Machine Learning for z/OS can assist in real time decision-making and automation staying in mainframes. The solution can be implemented in two steps:

1. Acceptance probability indication at the time of placing change request (Predictive Analytics) - Analyze the data to find acceptance patterns which can be used as input to make prediction and give an indication at the time of placing the order change on "how likely the manufacturing facility may accept/reject the particular change request"? This will help the personnel to decide if he should go for the change request.
2. Feed repetitive change requests as input directly to Forecasting/Ordering system - Analyze the data for potential repetitive change requests. These configurations can be back propagated as input to forecasting/ordering. Doing this has the potential to cut down change requests made. Dealer preference used at the time of vehicle allocation to dealers can also be considered at the time of placing firm orders.

Benefits:

- Significant reduction in the number of invalid requests and thereby reduce the overall workload which helps to process the valid requests faster.
- The insights on repetitive change requests helps to include them in the default package and thereby reduce the change request volume further down

Use Case #2: Using predictive analytics for efficient ATM cash flow management

Problem Statement:

Customers are heavily dependent on the Automated Teller Machines (ATM) for convenient and easy access to funds at any time of day or night. Inability to dispense cash at any of these ATMs due to cash shortage can result in customer dissatisfaction. Accurately forecasting the cash requirement at ATM is a challenge faced by all financial service providers.

If ATM is not loaded with enough cash, it gets emptied soon and can result in customer dissatisfaction. If ATM is loaded with more cash than required, it results in dead money for the bank which can invest somewhere else. If the ATM is loaded with incorrect denominations, customer unable to withdraw money of their required denomination despite ATM still having money.

The traditional 'Cash Flow Management' procedures focus on optimizing deterministic characteristics of the process and do not address the stochastic nature of the process arising out of fluctuations in cash demand. The traditional approach has been more intuitive relying on reactive measures.

The current trend is to adopt the scientific approach of predicting the cash requirement by studying historical data and applying machine learning. However, this requires more real-time access to monetary transactions of the customers which are highly confidential and managing data security is a primary concern

Challenges:

- All the transaction data resides on the mainframe. Offloading the data to data lakes/ data warehouses to do analytics can result in data latency, security risk, less efficient processing, significant cost and complex architecture

Proposed Solution:

Predict the optimal amount of cash replenishment for each ATMs using machine learning model. Use ARIMA method for Time Series Forecasting. Do the analytics and cash forecasting on mainframe itself leveraging IBM Open Data Analytics for z/OS which provide a secured real-time access to financial transactions.

The current model is not able to predict seasonal fluctuations in the cash requirements. Also, the current model does not have any way to predict the number of notes in different denominations that would be required at any particular ATM. To address these limitations, preemptive model based on 'Time Series Predictive Analytics' can be used to improve forecast accuracy.

The classical time series analysis procedures take into account the trend [a long-term monotonic change of the average level of the time series], the trade cycle [a long wave in the time series], the seasonal component [fluctuations in time series that recur during specific time periods] and the residual component [any other influences on the time series].

The below flow chart shows the various steps involved in the time series-based solution.

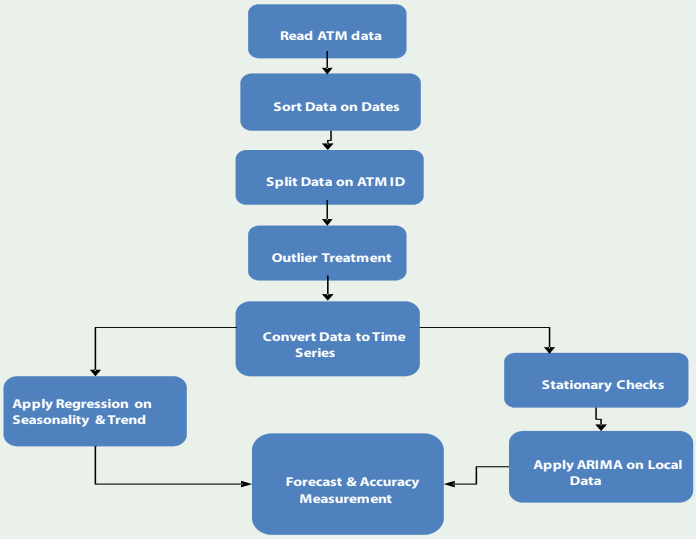

The model with ARIMA (1, 0, 0) was used which gave an accuracy up to 90%.

Benefits:

- More accurate prediction of ATM cash requirement as near real-time data is used for training and scoring.
- No data movement is required and there by avoid the data latency and security risk.
- Reduced cost and improved efficiency as training and scoring of models are done at data source

Application of Analytics and machine learning on mainframe in other industries

The two use cases depicted in this paper are related to Manufacturing and Financial Services. However, the real-time analytics and machine learning capabilities on mainframe is applicable to all the industries. Insurance, Card Payments and Retail industries have wide verity of uses cases that are more prevalent for adoption of IzODA platform. Few of the other interesting uses cases are fraud detection in card payments, identification of loan payment defaulters based on spend and payment pattern, early identification of client churn, identification of next best offer, inventory / stock anomaly detection, etc.

Conclusion

Assess your enterprise's data lake and data warehouse infrastructure for its flexibility, timeliness (a.k.a data relevance), and quality. Currently the data from mainframes may be ETLed at certain business-defined frequency and it is quite possible that the data feeding to the data lake will continue considering enterprise's holistic vision.

At the same time, assess the impact and the competitive edge you will get on performing real time analytics at the transaction sources directly (i.e. on the mainframe data directly). Ideally, accessing real time data is always better provided it is cost effective and does not impede the regular business operations and system availability.

IBM's products like IBM Open Data Analytics for z/OS as a comprehensive runtime and Watson Machine Learning for z/OS focused on end-to-end model life cycle management are powerful options to explore for medium and large mainframe enterprises if

- *The enterprise sees a positive impact to its business through real time decision making or overall timeline reduction in the business processes*
- *The enterprise has envisioned and positioned IBM mainframe as key infrastructure to run its critical business applications and modernize over the next decade or so*
- *The enterprise already has established open source products involving Python, Scala, Anaconda, and its supporting technology components part of its enterprise technology stack*

About the Authors

Prabhat Kumar Das
AVP, Delivery head of Manufacturing, Infosys

Prabhat is AVP and Delivery Head for Application Development and Maintenance Services for the Manufacturing and Resources verticals and Center Head for Infosys Bhubaneswar. He has more than 25 years of industry experience and has spent more than 15 years in Infosys. At Infosys, he has played a multitude of roles including leading delivery of large and complex transformational programs. Prabhat is also responsible for driving Corporate and Unit initiatives and anchoring digital offerings like Enterprise Agile & DevOps, Modernization etc. for global manufacturing clients.
Prabhat holds a Bachelor of Engineering degree in Computer Science & Engineering from the National Institute of Technology, Rourkela and has completed his Executive Leadership Program from Stanford Graduate School of Business, USA

Jaydip Sanyal
AVP, Senior Industry Principal, Mainframe Modernization Practice, Infosys

Jaydip Sanyal is an associate vice president and is part of the leadership team of the Modernization track in Infosys. Sanyal and his team are responsible for the strategy and technology behind modernizing legacy and proprietary technologies running on monolithic z/OS, AS/400, Java, .Net based architectures to new generation open source architectures on public and private cloud. He has participated in numerous modernization and open source events as a speaker and a panelist. He is an avid blogger on how you can modernize mainframes and make it digital ready. Sanyal has a degree in engineering and a postgraduate MBA degree from IIM with over 22+ years of experience in IT.

Suchindran Muthu
Principal – Enterprise Applications, Infosys

Suchindran is Principal – Enterprise Applications at Infosys focused on Legacy Modernization and cloud transformation solutions. His experiences involve major enhancements and transformation program involving mainframe integration, application transformation out of mainframes and modernization. He has 13+ years of IT engineering and Tech consulting experience and is a certified AWS Solutions Architect.

Nirja Parayane
Technology Architect, Infosys

Nirja is a Technology Architect at Infosys Ltd having 16 years of IT experience. Her primary skills are machine learning and big data technologies – Hadoop, HBase, Hive. She holds master's degree in data science and has implemented multiple machine learning use cases using Python, R, Spark etc. Apart from machine learning, she also has worked as architect for legacy modernization projects.

Tushar Tawale
Senior Technology Architect, Infosys

Tushar is a Senior Technology Architect at Infosys Ltd having 16 years of experience in analysis, architect, design, development and implementation of Enterprise Systems. His main expertise is as a Mainframe application architect, data architect and legacy modernization architect. He has exposure to MIPS optimization, service enablement, DevOps, re-engineering of mainframe applications. He also has experience on IzODA - Analytics platform on mainframe systems.

Rajaram VR
Senior Technology Architect, Infosys

Rajaram is a Senior Tech Architect , with wide spread technical, hands on expertise, across a repertoire of technologies. His core technology expertise is on open systems, and recently has expanded to include , key areas of focus, on modernization, legacy technologies, specializing in product space of AS400 re-hosting and auto code converters. He brings in extensive experience and design knowledge, specially, in the space of Database migration , and API approaches to data as a service, sourcing data from legacy platforms like Mainframe (hierarchical databases etc.). His recent interests are , large data compilation, transformation, no-SQL technologies.

Mythili Venkatakrishnan
Distinguished Engineer, IBM Z, Analytics Architecture & Technology Lead

Mythili is an IBM Distinguished Engineer and is the IBM Z Architecture and Technology lead as well as CTO for Z Financial Services Sector. Mythili has been with IBM for over 29 years, all in the IBM Z environment, working with clients in various capacities. Her focus areas have been diverse, but current responsibilities include leading the technical strategy and system structure definition for analytics and cognitive technologies such as Spark z/OS as well as helping clients to digitally transform their core systems leveraging their Z investments. Her previous roles have included industry solutions, business resilience architecture and design, availability management, systems design, and solution prototyping.

Nasser Ebrahim
Solution Architect, Data and AI on IBM Z

Nasser Ebrahim is a Solution Architect with IBM Systems Lab, focusing on Analytics and Cognitive solutions on IBM Z with Global System Integrators. Nasser has 22+ years of IT experience which includes system programming, mainframes technologies, Java, Spark, Kafka, Watson, Machine Learning, Deep Learning and Swift. He was Java Current Release Technical Leader with IBM Software Lab before taking the current role of solution architect.

David A. Rice
IzODA Chief Iteration Manager & Technical Lead of Scale Adoption

David Rice is the Chief Iteration Manager for the IBM Open Data Analytics for z/OS project and the technical lead of scale adoption. His primary focus is technical enablement of customers and business partners who want to use IzODA

For more information, contact askus@infosys.com

Infosys®
Navigate your next