



BATCH MODERNIZATION USING THE OPEN DATA ANALYTICS PRODUCT SET

Abstract

Application modernization has become critical for enterprises looking to embrace new opportunities. Batch jobs play an important role in enterprise IBM Z[®] workloads and, hence, require a custom approach for modernization. The right approach transforms batch applications while considering the end-to-end lifecycle of batch jobs.

This paper focuses on the different aspects of batch modernization like identifying the right modernization candidates, choosing the implementation language and more. The paper also elucidates batch modernization tools and solutions from IBM and Infosys that helps enterprises reduce cost and mitigate risk through efficient data processing.

Introduction

IBM Z mainframes are the backbone of many enterprises and plays a pivotal role in the daily operations across industries. By leveraging immense amounts of data and insights, IBM Z has transformed from being a mere platform that support operations to a revenue generating platform playing a central role in hybrid multi-cloud environments.

There is a common misunderstanding that the IBM Z and its technologies are aging and outdated. The reality is that latest technologies such as cloud, artificial intelligence (AI), blockchain, Internet-of-Things (IoT), and open-source packages can be leveraged on IBM Z. With new generation programming languages such as Java™, Python™, Node.js™, Swift®, etc., and open source based integrated development environments (IDEs), continuous integration/continuous delivery (CI/CD) pipelines are providing a similar developer experience as on any other platform. This developer experience is not limited to only new generation languages but also applies to legacy languages like COBOL, PL/I and Assembler.

IBM Z workloads can be broadly classified into online transactional processing (OLTP) and batch processing. For almost all enterprises across industries, batch processing is still a fundamental and mission-critical component. Applications on IBM Z are no different; with some enterprises running batch workloads up to 60% of total workload. Mission-critical batch windows often run on the IBM Z with minimal on-site staff (or even fully unattended in some cases), thereby processing large quantities of data at relatively low cost with great reliability.

Batch processing is used for a variety of applications. Some common examples are:

- Bulk updates to data/database (end-of-day transaction processing)
- Generating reports for daily, weekly, or monthly transactions
- Creating weekly or monthly account statements for customers of an organization
- Payroll processing of an organization
- Creating backups of files and databases for disaster recovery purposes
- Archiving historical data at periodic intervals
- Extract, transform, load (ETL) jobs to offload data to data lakes and data warehouses

Batch processing has significant advantages when executing repetitive logic. Static data is read only once, cached, and then used throughout the program. Thus, batch processing, i.e., bulk processing of massive amounts of transactions during non-office hours, remains a viable and strategic option.

Advantages of batch processing over OLTP

Batch processing is sometimes considered as the bulk processing of OLTP logic. However, there are some unique advantages of using batch jobs over OLTP:

- Real time processing is not always possible due to unavailability of dependencies. In such situations, batch processing is preferred
- In cases where extensive repetitive processing is needed, batch jobs are more efficient than OLTP due to their ability to cache static fields
- When bulk processing of data and online transactions happen simultaneously, bulk processing can negatively impact the service level agreements (SLAs) of online transactions and hinder user experience. This can be avoided by running bulk processing as batch jobs at later point of time
- Batch jobs ensures optimal system utilization during off peak hours
- Building periodic interface files with logical grouping of records that are to be sent to the information systems of partner organizations, it may be more efficient and reliable to send one big file with many records every day rather than each record in real-time



Strength of z/OS for batch processing

Compared to other platforms, IBM z/OS® provides high quality of service through its reliability, availability, security, scalability, and serviceability. In fact, IBM Z processes 68% of the world's production workload capacity with only 6.2% of server load demonstrates the superior efficiency of z/OS systems. It offers trusted computing with the highest security certification for commercial servers. The application availability of 99.99999% in IBM z15™ is the highest in the world.

The primary advantages of using z/OS for batch processing are:

- z/OS provides a dedicated batch processing environment with a job scheduler, job entry subsystem (JES), a spool and job management tool, and system display and search facility (SDSF)

- z/OS Workload Manager makes it possible to run many parallel batch jobs while sharing resources that are prioritized according to SLA specifications
- z/OS can provide maximum up-time through a unique clustering technology called Parallel Sysplex®
- Failures in batch programs can be recovered automatically. Execution of a batch program can be deferred to another logical partition (LPAR) or system in case of a disaster
- With z/OS Parallel Sysplex, peaks in batch processing can be scaled throughout the system by acquiring additional resources when needed
- z/OS ensures better overall utilization through resource sharing and data proximity

Challenges with batch processing

Currently, batch processes are optimally implemented in terms of performance, process control and monitoring. However, there are some challenges that can be addressed through batch modernization by adopting new tools and technologies. The main challenges with current batch processing methods are:

- **Programming languages:** Most of the applications on the IBM Z are still written in traditional procedural programming languages such as COBOL and PL/I. These languages cannot support new functional requirements such as creating PDF documents, sending e-mails, etc.
- **Availability of skills:** A rising concern for many companies is the difficulty in finding talent that is skilled in certain traditional technologies
- **Shorter batch windows:** The traditional batch window needs to be shortened to accommodate the increasing demand for OLTP windows. Business agility also demands that batches run quickly, thereby mandating shorter batch windows in order to achieve business outcomes faster
- **Running batches on-demand:** Besides enabling chunk processing of huge batch data, there is a need to run batches on-demand for business agility without affecting the SLAs of OLTP
- **Difficult to maintain:** The number and complexity of batch programs can result in long turnaround time to implement new business requirements, especially when batch programs contain numerous lines of code for formatting and transforming data

- **Application agility:** The time taken to introduce changes in modern/functional languages is shorter compared to legacy languages



Approach to batch modernization

As batch workloads are critical for organizations, batch modernization is an important aspect in the application modernization journey. Batch modernization needs an approach for transforming batch applications considering the end-to-end lifecycle of batch jobs.

Some important aspects to be considered for batch modernization are:

- Identifying the right candidates for batch modernization
- Determining the languages of implementation
- Discovering business rules and dependencies of existing applications
- Enabling simple and easy access to data from multiple data sources
- Enabling fast and efficient processing of data through parallelization
- Allowing easy consumption of job outputs by downstream applications
- Handling different types of batch workloads

These aspects of batch modernization can be achieved in z/OS using the following offerings from IBM and Infosys:

- IBM Z Platform for Apache Spark
- IBM Data Virtualization Manager
- IBM Application Discovery and Delivery Intelligence (ADDI)
- Infosys Live Enterprise Application Development Platform
 - Infosys SMF Log Analyzer
 - Infosys Business Rules Extractor



Journey to Open Data Analytics

Journey to Open Data Analytics for z/OS integrates key open-source analytics technologies with advanced data access and abstraction services. Designed to simplify data analysis, it combines open source run times and libraries with analysis of z/OS data at its source. This reduces data movement and increases the value of insights from current data.

This journey consists of the following components:

IBM Z Platform for Apache Spark:

This is built on Apache Spark™, a high-performance general execution engine for large-scale data processing. One of its key features is the ability to perform in-memory computing. Unlike traditional data processing technologies, Spark allows caching of intermediate results in-memory rather than writing them to a disk, thereby dramatically improving the performance of iterative processing.

IBM Data Virtualization Manager (DVM):

This component provides integration capabilities for IBM Z and other off-platform data sources. The data service can provide Spark or Python applications with optimized, virtualized, and parallel access to a wide variety of data.

Python AI Toolkit for z/OS:

Python AI Toolkit for IBM z/OS is a library of relevant open-source software to support today's artificial intelligence (AI) and machine learning (ML) workloads. It is a collection of Python packages for IBM Enterprise Python that can be installed and managed using Package Installer for Python (pip), the common Python package manager. These packages are provided to pip from an IBM-hosted PyPi-style repository, leveraging supply chain security, that makes your software management experience common across your Python environments.

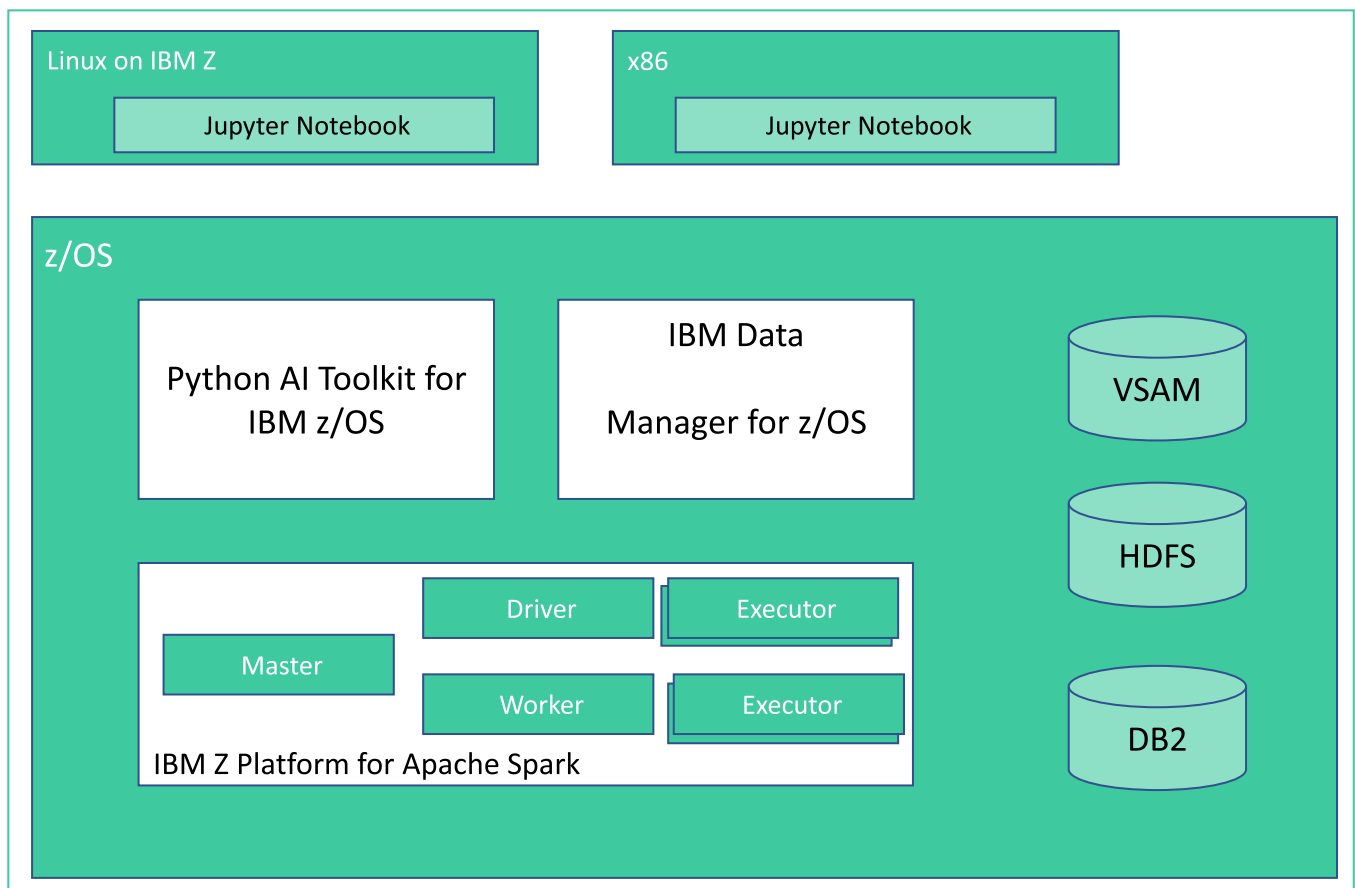


Fig 1: Journey to Open Data Analytics on z/OS

The products can also integrate with the IBM z/OS Connect Enterprise Edition if clients want to achieve consistent management across all RESTful services on z/OS.

Features and advantages of leveraging the products in the journey

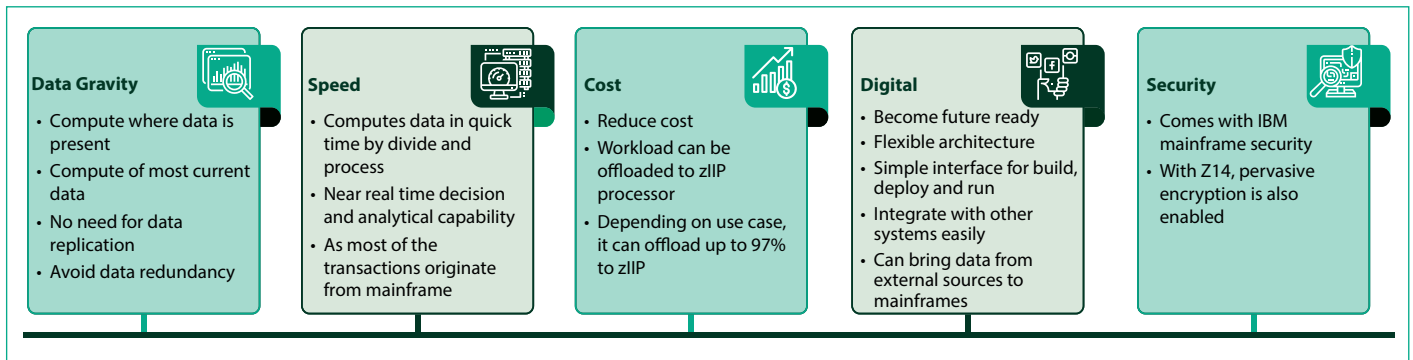


Fig 2: Features and benefits of leveraging the products in the journey

Some key benefits of Open Data Analytics:

- It eliminates the need to consolidate data before processing or analytics. It leverages abstractions of data from multiple disconnected sources on z/OS and other platforms and integrates them to perform a variety of large data processing using open-source technologies
- It provides industry standard open-source technologies and offers an optimized, in-memory approach. It delivers analytics results that can be integrated with business processes to discover previously hidden value
- It helps developers and data scientists quickly develop analytics applications using Java, Scala, or Python. It increases development capabilities with Apache Spark and Anaconda math, science, and parallel computing packages



IBM Application Discovery and Delivery Intelligence (ADDI)

IBM Application Discovery and Delivery Intelligence is the most comprehensive platform for IBM Z application understanding and modernization. It consists of two components:

- IBM Application Discovery for IBM Z – It helps developers better understand their application landscape through detailed graphical views, reports, flow diagrams, and many other features
- IBM Application Delivery Intelligence for IBM Z – It is a web-based tool that provides insights of the applications analyzed by

IBM Application Discovery for IBM Z. It uses a dashboard to present application details like run-time performance, static code quality and testing information, thereby enabling early problem detection. It also allows users to discover business rules that are embedded in the logic of web, online and batch applications. These rules can be identified, validated, and transformed. Dashboard information can also be exported through APIs to be used by other tools.

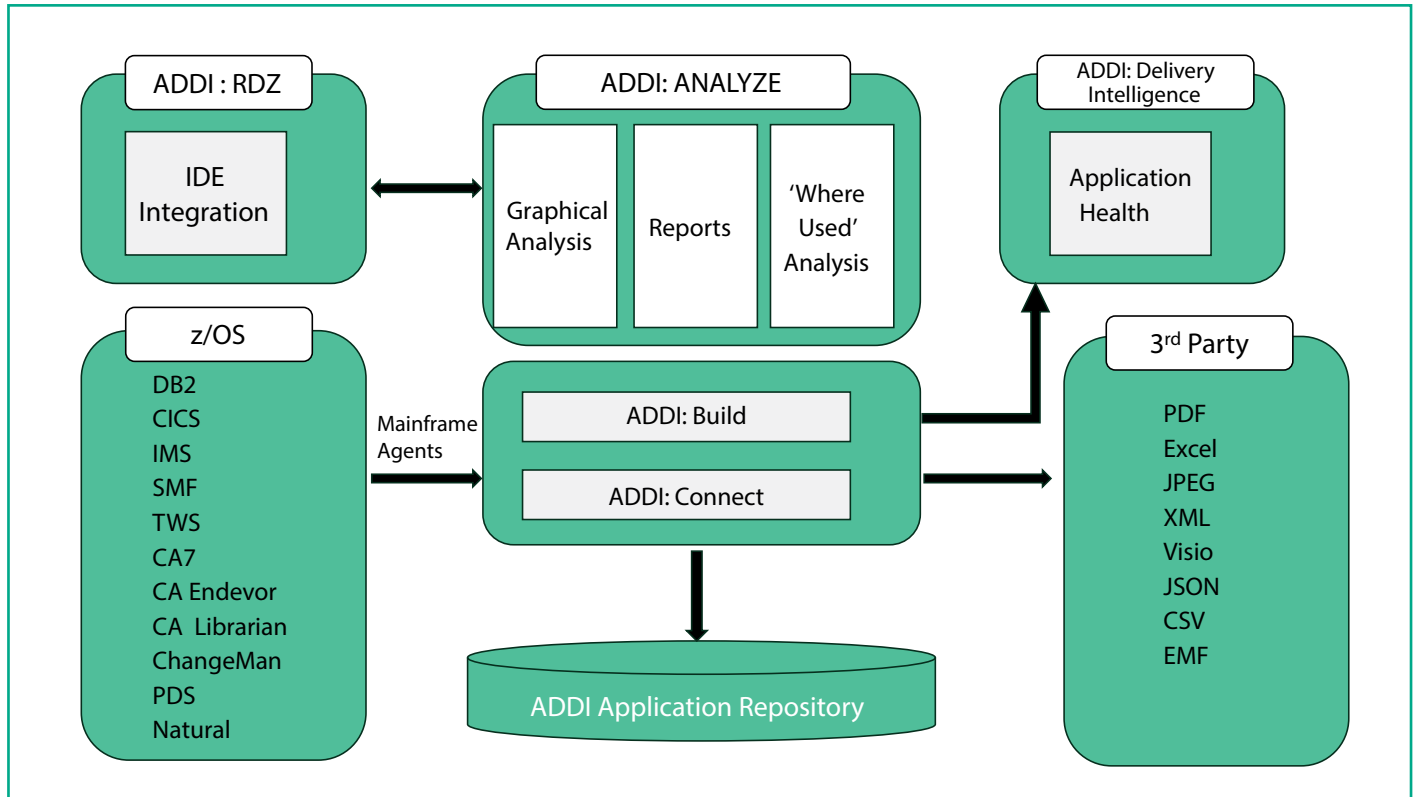


Fig 3: Architecture of ADDI

The output of ADDI processing in the ADDI application repository can be fed into DVM. The data studio associated with DVM provides a wizard-like interface for clients to choose what data environments associated with the code artifacts should be

virtualized. With this approach, ADDI can be used to identify the landscape of batch applications and dependencies while DVM can be used to identify and virtualize the right subset of data artifacts in preparation for batch modernization.



Identify the right candidate for batch modernization

Since all batch workloads may not be suitable for modernization, it is important to prioritize these based on certain criteria, thereby ensuring maximum return on investment. Some of the criteria include long-running jobs (or job streams that take a lot of time) and high CPU consuming jobs or job streams.

The above-mentioned tools and processes can help identify the right candidates for modernization based on multiple parameters that are extracted by static code analysis as well as interaction with SMEs for a top-down approach.

Batch modernization begins with a detailed assessment. Infosys uses a combination of function-driven and tool-based analysis enabled by IBM and Infosys in-house tools. Infosys also conducts meetings and workshops and leverages proven expertise and knowledge to understand the business requirements.

In addition to IBM ADDI, Infosys Mainframe Dynamic Attribute Extraction tool is useful for assessing applications. This tool can generate metrics at a capability/use case level that helps determine whether the workload is a right fit for modernization. Infosys SMF Log Analyzer is another tool that identifies the active inventory, thereby helping organizations save costs on modernizing entire applications. The tool can identify redundant, long-running, step-level long running, and high CPU consuming jobs in an application. It can also measure the run frequency of a job for a particular duration.

Determine the language of implementation

Most batch applications on the IBM Z were written many decades ago in traditional procedural programming languages such as COBOL or PL/I. While such jobs may run at optimal performance, control, and monitoring, they may be insufficient to meet the current business challenges and needs of clients.

New functional requirements such as creating PDF documents, sending e-mails, etc., are difficult to implement with legacy

languages like COBOL. Moreover, many organizations want to reduce the batch window. When dealing with large volumes of enterprise data, this is only possible by increasing the parallel processing of bulk data – and popular frameworks for parallel processing of big data support only new generation languages. Finally, addressing other major challenges with batch jobs (like availability of skills on legacy languages and maintaining huge monolithic batch applications) is possible only when the applications are refactored or rewritten.

The most popular big data processing engine is Spark, which supports languages like Java, Scala, Python, and R. Hence, it is important to choose one of these languages for batch modernization. R is primarily used by data scientists for predictive or descriptive analytics and is not ideal for batch jobs. Thus, we will consider only Java, Scala, or Python for batch modernization.

Java and Scala are JVM-based languages that produce similar bytecodes and run on the same Java virtual machines. In the earlier days of Scala, the primary differentiator was the functional programming paradigm. This difference was significantly reduced with the introduction of lambdas in Java 8 along with further enhancements. Scala is more concise whereas Java is more verbose in programming; but Scala is harder to learn compared to Java. Java is considered to be more readable with more availability of marketplace skills. Bearing this in mind, it is up to the programmer to choose either Java or Scala since both the languages use the same runtime and have similar performance.

When comparing with Python, Scala has better performance. It supports powerful concurrency through primitives like Akka actors. In Python, the Global Interpreter Lock often causes issues when maximizing parallelization during data ingestion and execution. Moreover, Scala is type safe and, hence, less error prone during runtime. However, Python is easy to learn and less verbose compared to Scala.

In summary, Python is ideal for proofs of concepts (PoC), demos or small applications whereas Java and Scala are recommended for production applications including batch.



Discover business rules and dependencies using IBM ADDI and Infosys Business Rules Extractor

For batch modernization, IBM Application Discovery provides detailed reports on the inventory of batch applications, batch flow visualization and code complexities. Users can also access information about scheduler and dead code. This can help identify components that need not be modernized. Business Rule Discovery of IBM Application Delivery intelligence helps unravel potential business rules within the artifacts identified by Application Discovery. Application Delivery Intelligence scans the files and automatically discovers keywords and their implementation names.

Infosys Mainframe Modernization architects analyze the reports and leverage their domain knowledge to create the business rules document. This document is then reviewed with application SMEs to ensure that the functionality aligns with the technology. It is the basis for enterprise architects to easily modernize batches.

Infosys Business Rules Extractor tool can further aid this process. Besides the rules that are generated by ADDI, if additional documentation of rules is needed, Infosys Business Rules Extractor can centrally document all the rules with program management features like:

- Assigning tasks
- Updating completed status of tasks
- Tracking overall progress in terms of use cases lines of code (LOC) documented
- Creating multiple personas for different access levels
- Enabling rule chaining across modules based on documentation by analysts

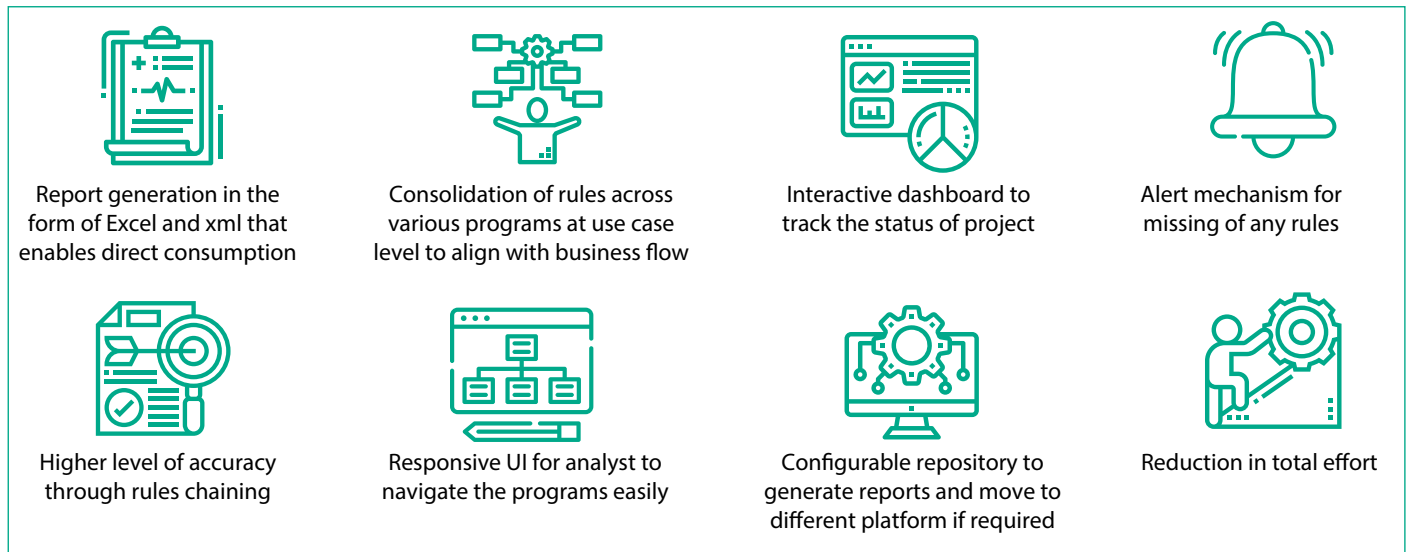


Fig 4: How Infosys Business Rules Extractor accelerates business rules extraction



Streamlining access to multiple data sources with IBM Data Virtualization Manager

The first step in data processing is accessing the data sources. Hence, for batch modernization, it is important to ensure simple and easy access to data sources. IBM Z systems comprise of many data sources like Db2®, Virtual Storage Access Method (VSAM), IMS™, etc., and each of these have different access methods. For example, Db2 is a relational database management system (RDBMS) and can be read using SQL queries. IMS is a hierarchical database and is accessed as parent/child records whereas VSAM can be read as a flat file using its schema. Hence, it is challenging for application developers to create and maintain batch applications with multiple data sources.

IBM Data Virtualization Manager is a simple, easy, and efficient way to access different data sources in IBM Z. These sources include Db2, IMS, VSAM, physical sequential dataset (PS), partitioned data set extended (PDSE), Adabas, Integrated Database Management System (IDMS), CICS® queues, virtual tape, System Managed Facility (SMF) records, Syslog, etc. Access is enabled through virtual integration into a single, logical data source that can then be imported into a Spark DataFrame for programming online and batch applications. The DVM also provides data integration facilities for off-platform data sources like Db2 for Linux®, UNIX® and Microsoft® Windows® (Db2 LUW), Oracle Enterprise, Teradata, Hadoop Distributed File System (HDFS), etc., so that the application can access distributed data sources while processing IBM Z data.

The main advantages of using the DVM for data access are:

- All data sources (including non-relational databases like IMS, VSAM, PDSE, etc.) can be accessed as relational data through virtual tables and views
- The DVM is highly IBM z Integrated Information Processor (zIIP)

eligible, thereby minimizing total million service units (MSU) consumption for accessing data

- DVM performs in-memory caching of data that is being read from data sources until the application is ready to consume it, thereby accelerating data access
- DVM has a unique 'direct-read' feature for Db2 and IMS that performs a bulk fetch of data directly from the back-end datasets. This streamlines data access, avoids negative impact on the service level agreements (SLA) of other applications that access IMS and Db2 subsystems and increases overall zIIP eligibility
- The DVM is fully integrated with z/OS prime features like workload management and security and hardware features like Simultaneous Multithreading (SMT2), IBM zEnterprise® Data Compression (zEDC), etc.
- DVM provides configurable data privacy through column or row obfuscation/removal across any data source
- DVM automatically generates optimized data access codes for faster application development

The DVM offers a client application/user interface called the Data Service Studio. The first step in accessing data is to configure the data sources (that were identified in the discovery phase) into virtual tables and views in the DVM using the Data Service Studio. Following this, all data sources can be seamlessly accessed from the Spark application using SQL statements without being concerned about the type of database or access methods of the original data source.

The DVM can pull the Db2 table information from the Db2 catalog and IMS information from the IMS RESLIB. The schema of VSAM datasets or flat files can be taken from the COBOL copybook or PL/I Include files.



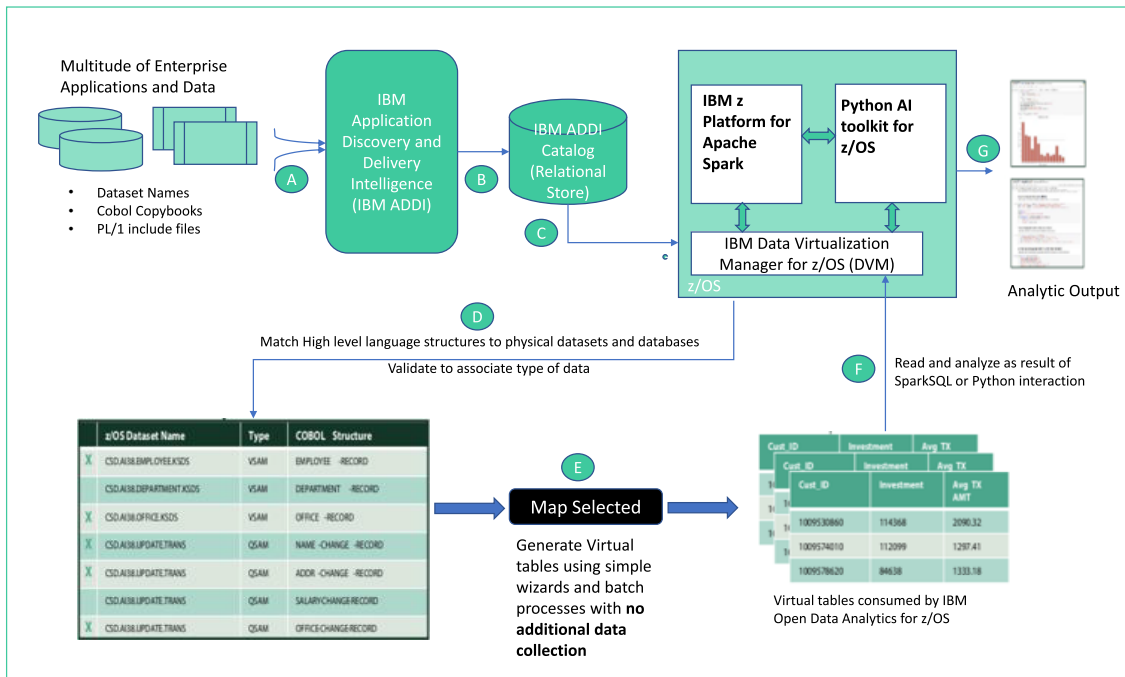


Fig 5: Using IBM ADDI to automatically discover/configure data sources to the DVM

IBM ADDI can be used to automatically discover and configure datasets and databases into virtual tables and views. Once batch applications and various data sources are fed to IBM ADDI, it uncovers the meta data of data sources and feeds this into the DVM. It also matches the high-level language structure to physical datasets and databases. The DVM then validates this according to the type of data. Virtual tables and views are then generated using simple wizards without the need to collect additional data.

Efficient batch processing using Spark Runtime

Apache Spark is a fast and general-purpose cluster computing framework. It has an in-memory data processing engine for big data processing along with in-built modules for streaming, SQL, machine learning, and graph processing. Apache Spark is well-known for its speed, ease of use and extensibility. When used on z/OS with zEDC, it optimizes compression.

Resilient distributed datasets (RDD) are the fundamental data structure in Apache Spark. These are an unchanging distributed collection of objects. Each RDD is divided into logical partitions that can be computed on different nodes of a cluster.

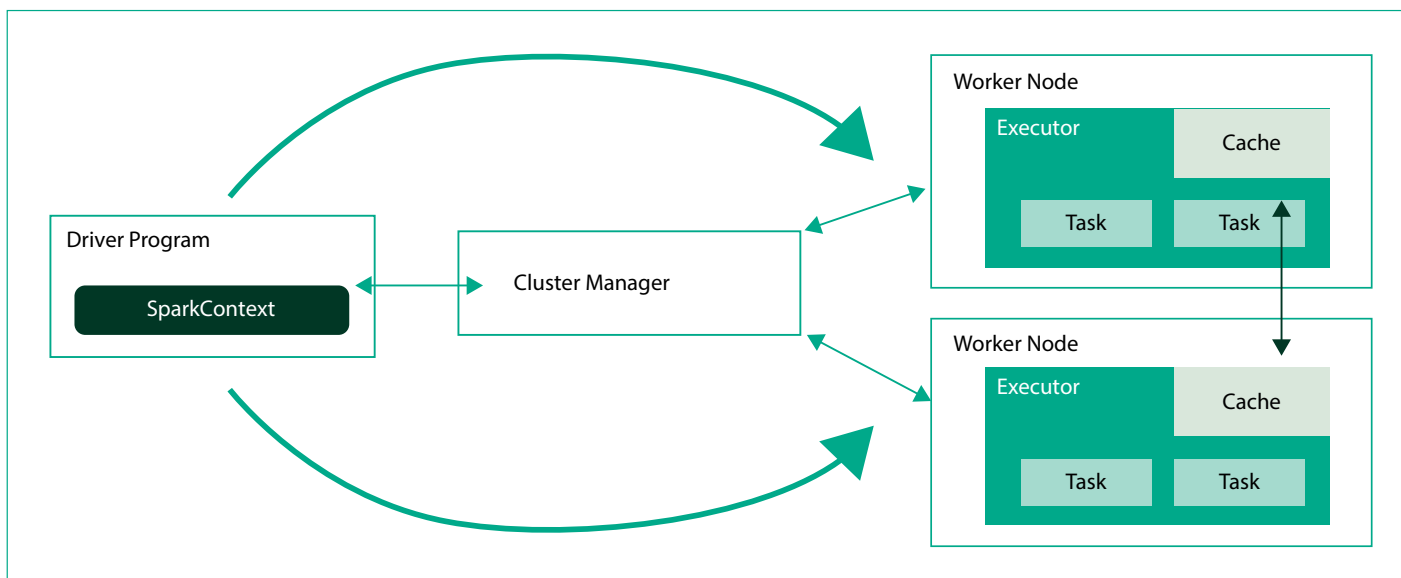


Fig 6: Components of the Spark cluster on z/OS

The Spark cluster on z/OS consists of the following components:

- **Driver program** – Spark Driver is the program that declares the transformations and actions on RDDs and submits requests to the master. It runs the main function of the application that prepares the SparkContext. The SparkContext object in the main program orchestrates all the activities within a Spark cluster. Each driver schedules its own tasks.
- **Cluster manager** – Cluster manager is an external service for acquiring resources on the cluster. The SparkContext object connects to the cluster manager, which allocates resources across applications. Spark applications run as independent sets of processes on the Spark cluster.
- **Worker node and executor** – Any node that can run application code in the cluster is a worker node. An executor is a process launched for an application on a worker node that runs tasks and keeps data in memory or Direct Access Storage Device

(DASD) storage. Each application has its own executor processes, which are active for the duration of the whole application while running tasks in multiple threads.

When submitted as a Spark job, the batch application distributes the work to multiple executors for parallel processing, enabling efficient data processing in lesser time compared to traditional batch jobs.

Here is an example to explain how batch processing using IBM Z Platform for Apache Spark runtime improves efficiency.

Assume an enterprise has two systems of records (SORs) containing billions of rows of data. The business demands that these two SORs are synced on a weekly basis. In the traditional batch method, this entails having a series of jobs that read data in a serial manner. The data is then compared using utilities like 'sort' along with data manipulation, data convergence and divergence. Executing this entire series of jobs is time-consuming and expensive.

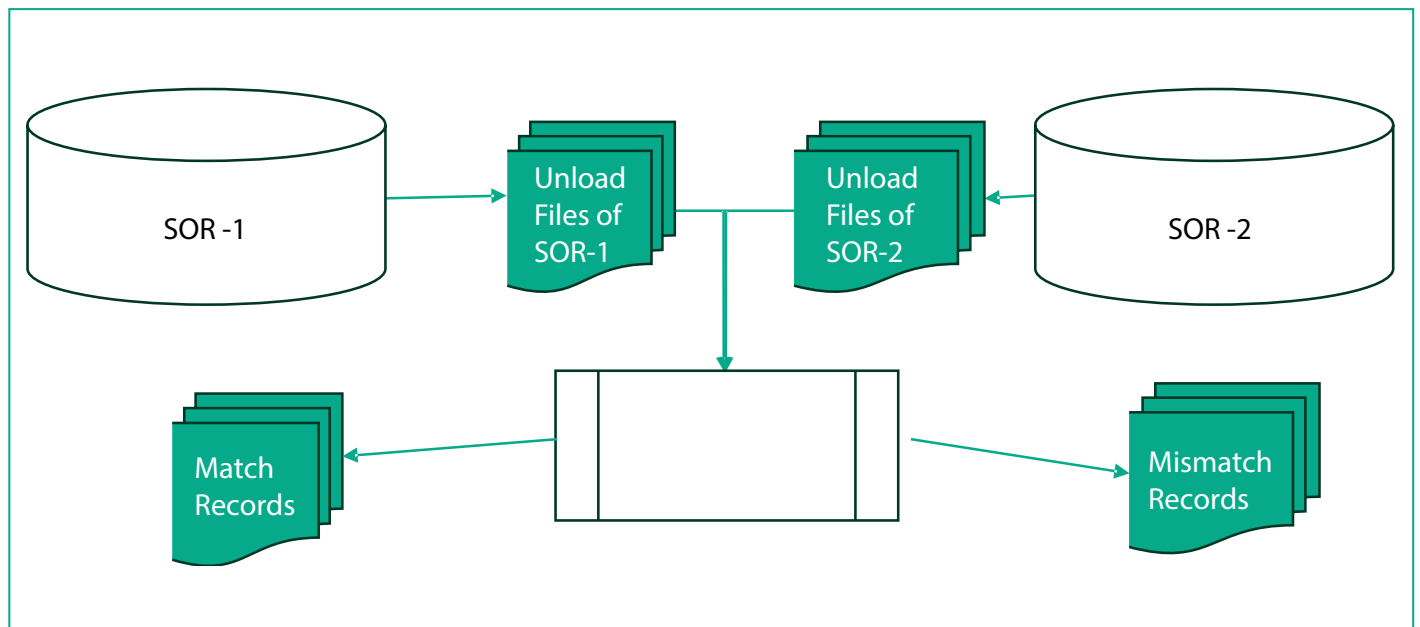


Fig 7: Traditional batch jobs to sync two system of Records

The same job can be processed very efficiently, cost effective and less time consuming in a Spark environment. The spark job gets split into independent smaller tasks and distributed

among the executors in the worker nodes for parallel processing. Spark would keep all the data in memory, making all computes very efficiently, cost effective and less time consuming.



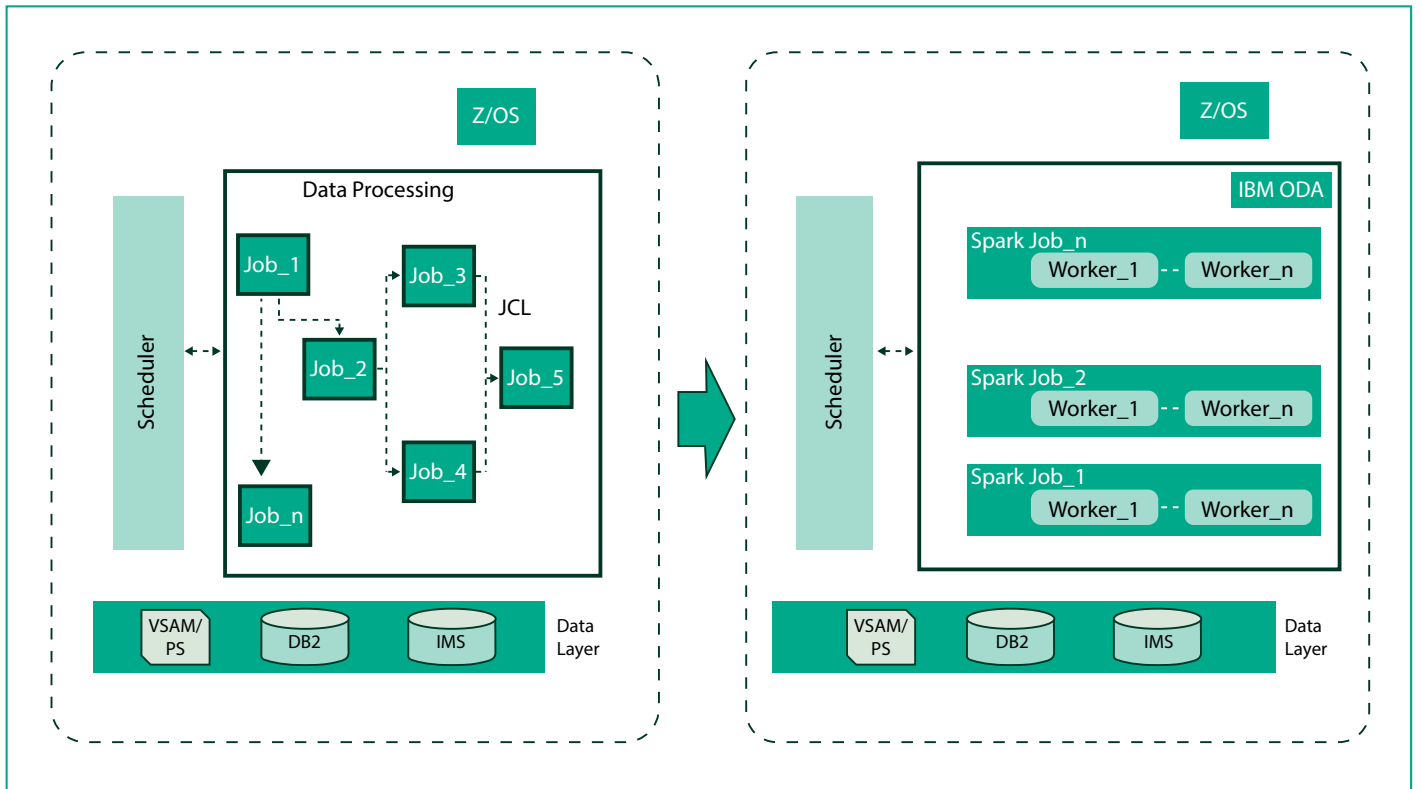


Fig 8: Transforming a traditional batch job to parallel processing in Spark environment

Easy consumption of batch output using Apache Spark result data store

The Spark job output is stored in an in-memory virtual table called the 'result data store'. The output can be accessed from anywhere through REST APIs, JDBC, ODBC, or MongoDB interface using the DVM. If the output has to be stored permanently, it can be pushed to a Db2 table as a result data store, which can be accessed like any other data source.

If the batch output is fed as an input to another batch job, then the output of the previous batch job need not be persisted and can be processed as in-memory storage. However, if the batch job output has to be accessed on-demand by other systems or applications, then it is better to persist the data.

Handling different types of batch workloads

During the batch modernization journey, various types of batch workloads can be handled differently to improve efficiency and ROI.

ETL jobs form a significant batch workload when moving IBM Z data into data lakes or data warehouses for analytics or insights. An IBM analysis reveals that, for most enterprises, an average of 16% to 18% of total MSU consumption is caused by ETL jobs. The traditional approach of processing ETL jobs can lead to high data latency and significant processing cost. Moreover, moving data

and maintaining multiple copies can cause data inconsistencies, resulting in greater security risk and non-compliance. As ETL jobs typically run parallel to other batch workloads, it puts undue load on resources, significantly impacting the processing efficiency and SLA of all processing workloads.

IBM Z systems are enabled with seamless data processing, efficient data access and open industry leading frameworks like Apache Spark to run batch jobs efficiently while performing analytics and machine learning functions. Any risk arising from data movement can be averted or significantly reduced if analytics and data processing are performed at the point where data is generated.

Even when data needs to be accessed by other systems or shared with enterprise partners, its physical movement can be avoided. Data can be accessed on-demand through REST APIs or JDBC using the DVM. Preparation, distillation, aggregation, etc., of raw data can be done at the data source and the processed data or insights can be stored in the result data store. This allows other systems or business partners to access the processed data or insights on-demand from anywhere and at any time.

For batch jobs with read only data, the Db2-direct and IMS-direct feature to directly read data from backend VSAM datasets instead of going through the database manager is useful for IMS and Db2 databases. In this way, batch jobs can be processed along with online transactions without impacting the SLA.

Conclusion

For nearly every enterprise, batch processing is still a fundamental and mission-critical component with 40–60% of an enterprise's total workload running through batch. To effectively manage dynamic business requirements, stiff competition, dwindling resources, and technological disruption, it is important for enterprises to adopt batch modernization by leveraging new capabilities of IBM Z.

Embarking on batch modernization requires an approach that transforms batch applications while considering the end-to-end lifecycle of batch jobs. Using IBM Z Platform for Apache Spark, IBM Data Virtualization Manager for z/OS (DVM), and IBM Application Discovery and Delivery Intelligence (ADDI) along with Infosys solutions deliver specific capabilities for batch modernization. The business rules of the existing legacy batch application and its dependencies can be discovered with ADDI and Infosys tools. Access to different data sources, efficient processing of batch application and easy consumption of job outputs can be

achieved using the open data analytics product set analytics runtime on IBM Z.

The open data analytics product set combines efficient, fast, and cost-effective data access using IBM Data Virtualization Manager for z/OS and highly parallelize efficient batch compute with Apache Spark, leveraging in-memory features for intermediary data sets. This delivers a key advantage when used on the IBM Z where the vast majority of data for batch processing originates.

The traditional approach of ETL jobs and data movement leads to high costs, risk, and data latency. These overheads can be significantly reduced if analytics and data processing is done on IBM Z using IBM Z Platform for Apache Spark. If data needs to be accessed by other systems or shared with partners for business needs, IBM Data Virtualization Manager for z/OS can prepare the data at the data source and share only the processed data. This eliminates the need to move huge volumes of raw data, thereby improving efficiency and reducing cost.

Appendix

Acronyms:

ADDI	-	Application Discovery and Delivery Intelligence	ODBC	-	Open Database Connectivity
AI	-	Artificial Intelligence	OLTP	-	Online Transaction Processing
CI	-	Continuous Integration	PDSE	-	Partitioned Data Set Extended
CICS	-	Customer Information Control System	PL/I	-	Program Language One
CD	-	Continuous Delivery	POC	-	Proof of Concept
COBOL	-	Common Business Oriented Language	PS	-	Physical Sequential
CPU	-	Central Processing Unit	RDD	-	Resilient distributed datasets
DASD	-	Direct Access Storage Device	RDBMS	-	Relational Database Management System
DVM	-	Data Virtualization Manager	REST	-	Representational State Transfer
ETL	-	Extract, Transform, Load	ROI	-	Return of Investment
HDFS	-	Hadoop Distributed File System	SDSF	-	System Display and Search Facility
IDE	-	Integrated Development Environment	SLA	-	Service Level Agreement
IDMS	-	Integrated Database Management System	SME	-	Subject Matter Expert
IMS	-	Information Management System	SMF	-	System Management Facility
IOT	-	Internet of Things	SMT	-	Simultaneous Multithreading
JDBC	-	Java Database Connectivity	SOR	-	System of Record
JES	-	Job Entry Subsystem	SQL	-	Structured Query Language
LOC	-	Line of Code	VSAM	-	Virtual Storage Access Method
LPAR	-	Logical Partition	zEDC	-	zEnterprise Data Compression
LUW	-	Linux UNIX Windows	zIIP	-	z Integrated Information Processor
MSU	-	Million Service Units			

About the authors

Dr. Sameer Goel

Global Practice Head for Mainframe Modernization in Infosys and has been leading modernization programs since last five years. He is a Profit center head with 28+ years of experience of managing large IT global portfolio, large complex programs, stakeholder management and delivery excellence of the IT transformation projects. He also conceptualizes Go-To-Market Plans, Alliance Plans and created new opportunities through relationship building with business sponsors.

Mythili Venkatakrishnan

Distinguished Engineer with IBM Z and has been with IBM for over 30 years. Mythili is the IBM Z Financial Services Sector CTO and works with enterprise clients on digitally transforming and modernizing core systems for agile interaction with hybrid cloud. In that role, she collaborates across IBM as well as industry and ecosystem partners to enable clients to realize faster time to value via optimized integration of their IBM Z investments. In addition, Mythili is the technical lead for the IBM Z Digital Integration Hub focused on real-time information sharing between z/OS core systems of record and hybrid cloud. Her previous roles have included leading the analytics architecture and technology for IBM Z, business resilience architecture and design, systems design and solution prototyping.

Nasser Ebrahim

Technical Lead for IBM Z Digital Integration Hub and a Senior Solution Architect on Data and AI with IBM Systems Lab. As the technical leader, he works closely with clients and business partners to integrate IBM Z with Hybrid cloud for real-time information flow. Working with global systems integrators, he also specializes in analytics and cognitive solutions on IBM Z. He has over 25 years of IT experience, which includes systems programming, Java, Spark, Kafka, IBM Watson®, Machine Learning, Deep Learning, and Swift. He was Java Current Release Technical Leader with IBM Software Lab before taking on his current role.

Anjali Abraham

IBM Z DevOps Solution Architect with IBM Systems Lab. She has more than 17 years of IT experience, all on IBM Z. This includes application programming and solution architecting.

References

<https://www.infosys.com/services/application-modernization/offerings/mainframe-modernization.html>

<https://www.ibm.com/in-en/marketplace/open-data-analytics-for-zos>

https://www.ibm.com/support/knowledgecenter/en/SS3H8V_1.1.0/com.ibm.izoda.v1r1.izodalp/izoda.htm

<https://www.ibm.com/us-en/marketplace/app-discovery-and-delivery-intelligence>

<https://www.ibm.com/support/pages/ibm-application-discovery-and-delivery-intelligence-ibm-z-library>

<http://www.redbooks.ibm.com/abstracts/redp5217.html?Open>

<http://www.redbooks.ibm.com/abstracts/sg247779.html?Open>

<http://www.redbooks.ibm.com/abstracts/sg248325.html?Open>

For more information, contact askus@infosys.com



© 2023 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.