IBM Platinum Partner | Capgemini

MAINFRAME MODERNIZATION PATTERNS FOR FINANCIAL SERVICES

It's our observation that financial services organizations have been pioneers in adopting technology to help reduce operating costs, bring innovations faster and provide better customer experience. For many, IBM Mainframes have been part of their core IT platform for more than 5 decades due to powerful processing power, reliability and security. Even today, 85 of the top 100 banks use mainframes to run their mission critical IT workloads.

The new competition from FinTechs and BigTechs, evolving business models and regulatory and compliance risks are driving changes in financial services industry's technology needs. They are looking for a modern platform which can provide agility, security and scalability at a lower cost. They are investing in cloud computing and have started moving their non-critical workloads to cloud. Some of the organizations are already underway moving critical user facing applications and data to cloud to provide faster time to market, better user experience and improved scalability without comprising on security. However, some may still be reluctant to migrate mission-critical mainframe workloads to cloud due to risks of migration failure, very long migration timelines and little or no know-how of the old processes which may make the migration difficult.

Some of the challenges that organizations may face with the mainframe platform motivating them to look for alternatives are:

- **Shortage of skilled resources** - The existing mainframe skilled staff is retiring and is not being replaced by new talent due to lack of interest by younger generation engineers to learn mainframe skills
- **Outdated tools and development environments** – Many mainframe developers still use ISPF green screens to write, debug and test code, elongating the overall development and product release cycles, which could be a barrier for onboarding new staff who are used to work with modern development environments
- **Manual Dev and Operations processes** – Low penetration of DevOps practices for mainframe application development, causing two-speed IT where distributed development platform is more agile than the mainframe
- **Limited access to application and data causing interoperability issues** - The business logic hidden in legacy programs is not easily accessible to distributed applications, making it difficult for mainframe to operate in modern IT ecosystem
- **Inability to unlock the value of massive business data stored on mainframes** – The mainframe data stored on disparate database management systems is not accessible in real-time for analytics and machine learning processes
- **High cost of operation** – The technical debt accumulated over decades and unoptimized processes are contributing to higher costs of operations on mainframe platform

Several of these challenges are perceived rather than real. Many of the financial services organizations have not refreshed their mainframe tooling and technology stack over a decade and continue to run it as a legacy platform. With the release of z13 in 2015, the mainframe platform was redesigned to bring many innovations and modern tooling. IBM has been continuously investing in modernizing this platform with the release of z14 in 2017, z15 in 2019 and z16 in 2022. With each version, the modern mainframe hardware and software was designed to provide enhanced security, speed, resiliency and modern tools to interoperate in a Hybrid Cloud ecosystem.

The Hybrid Cloud infrastructure is a combination of one or more public cloud, private cloud and on-premises infrastructure. The hybrid cloud has the potential to facilitate faster code development and reduce time to market due to automation, consistent use of tooling and DevOps practices. We find that most organizations are embracing hybrid cloud strategy as their preferred cloud deployment model. The mainframe platform will continue to be an important part of this strategy and organizations need to start thinking long term on what should move to cloud and what can stay and be modernized on mainframe.

Capgemini and IBM help financial services customers navigate through their mainframe modernization journey. Capgemini provides a balanced approach towards mainframe modernization by choosing the right workloads to migrate to public/private cloud along with identifying in-place modernization patterns to modernize mainframe workloads.

The Capgemini In-place mainframe modernization solution focuses on modernizing developer experience, bringing agility and speed of delivery, providing real-time access to application business logic and data, and bringing AIML and data analytics closer to source data. Capgemini has created a mainframe modernization Experience Zone where customers can experience these modernization use cases in action and co-develop their use cases with Capgemini mainframe experts before implementing it in their environment.

BELOW IS A BRIEF SUMMARY OF EACH OF THESE USE CASES.

Modernizing developer experience

Mainframe developers have not seen significant changes to their development environments in the last few decades and they are largely based on 3270 emulator screens using the ISPF editor. Customers can embrace the modern integrated development environments (IDEs) of their choice based on industry preferred platforms such as Eclipse or VSCode to integrate with the mainframe. These are well known platforms used for developing applications written in newer generation languages and now the same platform can be used to develop applications on mainframe in COBOL, PL1, JCL and many other supported languages.

These are feature-rich IDEs and provide powerful tools to code, debug and test applications which are not available on traditional 3270 based mainframe development environments.

Implementing modern IDEs brings the experience of developing mainframe applications closer to distributed application development. It is easier to find new talent who have worked on these environments and can easily learn COBOL to support mainframe applications.

Figure 1. Side-by-side view of two IDEs: ISPF editor v/s Eclipse based IBM Developer for z/OS (IDz)

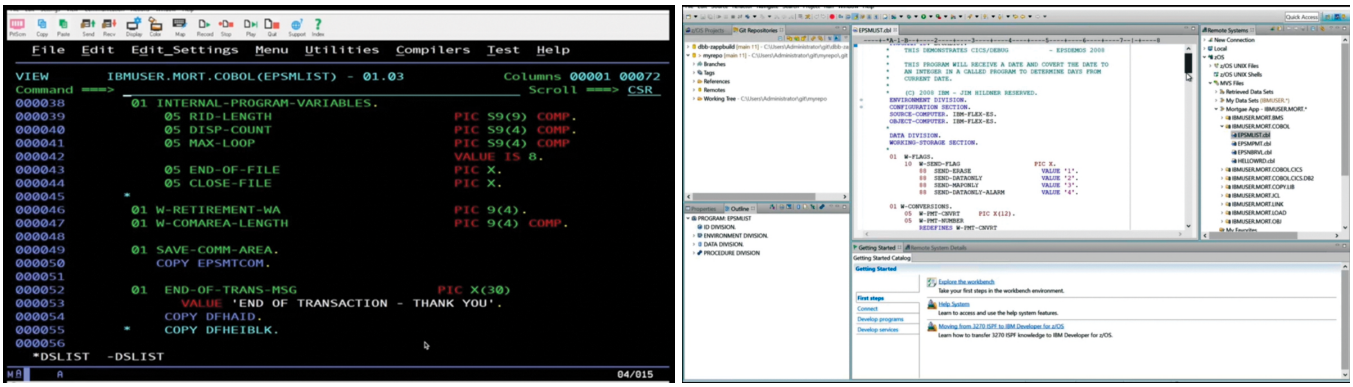

The success of this use case relies essentially on adoption of these tools by mainframe development staff. Many of the mainframe developers are using ISPF editors for decades and onboarding them to new IDEs will require training and coaching followed by regular feedback sessions. Once they see the power of these editors and efficiency gain they bring in their day-to-day work, they will adopt quickly.

Agility and speed of delivery using DevOps tools and practices

Organizations are adopting agile software development life cycle to deliver more functionalities at a faster rate. DevOps has been at the heart of this strategy and it has helped to speed-up the delivery cycles by shifting the role of delivery and operations from multiple teams to software engineers, providing them tools to automate the entire lifecycle using CI/CD toolchain.

Many organizations today either already have matured DevOps practice or in-process of implementing one, but the majority of these implementation only focus on distributed and cloud native technologies leaving the mainframe as an island to operate in the old world. This creates two-speed IT situation where one team is delivering faster than other.

Many of the open-source CI/CD toolchains can now be configured to work with the modern mainframe such as GitHub and Jenkins. Organizations can leverage these tools to create common DevOps framework to orchestrate the development, integration and deployment of an application to multiple target platforms including cloud and mainframe.

Figure 2. Sample reference architecture to implement DevOps pipeline for mainframe technology stack

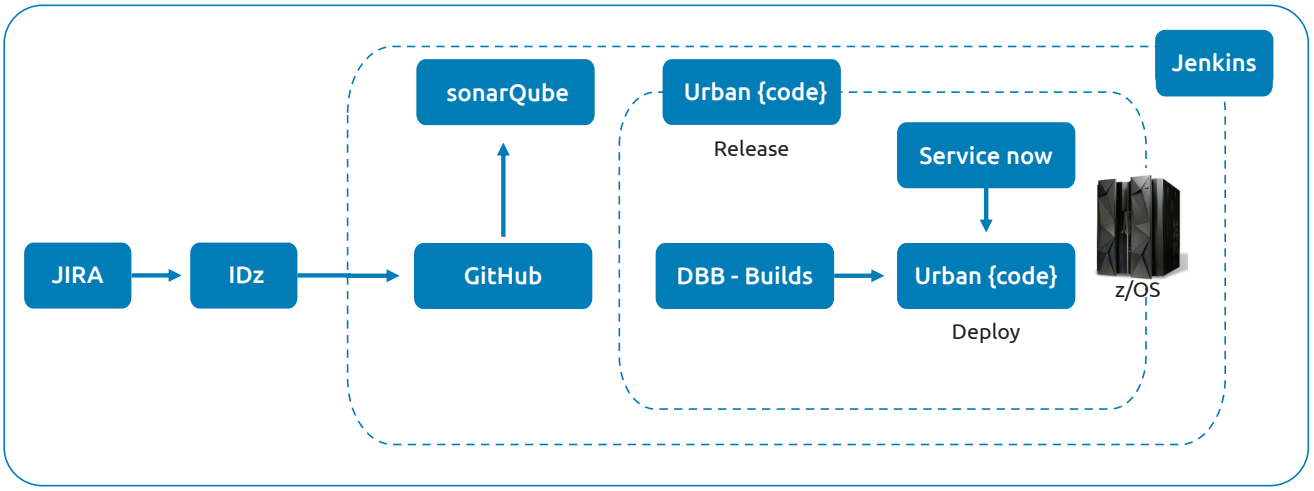

Activities	Requirement Management	Development Modernization	Source Control Management	Code Quality	Automated Unit Test	Build	Deploy	Release Management	Continuous Monitoring & Reporting	CICD	ITSM Integration
Tools Used	JIRA	IDz	GitHub	SonarQube	zUnit	Dependency Based Build (DBB)	Urban Code Deploy (UCD)	Urban Code Release	Omegamon / Instana	Jenkins	ServiceNow

Like the developer experience use case, the adoption is a key to success of any DevOps implementation. Start with the proof of concept to ensure that your setup works as expected, and then onboard the applications to DevOps pipeline based on the maturity of an application and development team.

Real-time access to mainframe applications and data

Digital transformation requires access to core applications and data to operate on omnichannel platform. Majority of financial services organizations have their core business logic on the mainframe written in old programs with little or no documentation. Re-engineering and rewriting these programs on a different platform is effort-intensive and error prone.

This modernization use case transforms mainframe COBOL programs to RESTful APIs which can be accessed by internal or external applications using open API standards. This not only ensures the seamless exchange of information across hybrid cloud ecosystem but also creates ways to monetize the organization's valuable intellectual property buried in these legacy programs.

IBM's z/OS Connect accelerates the creation of these APIs by providing easy to use API designer for developers to define the request and response parameters and create APIs in minutes. These APIs can be registered and discovered from an enterprise API management platform by any of the cloud applications running on public or private cloud.

Figure 3. z/OS connect EE for transforming mainframe assets to business APIs

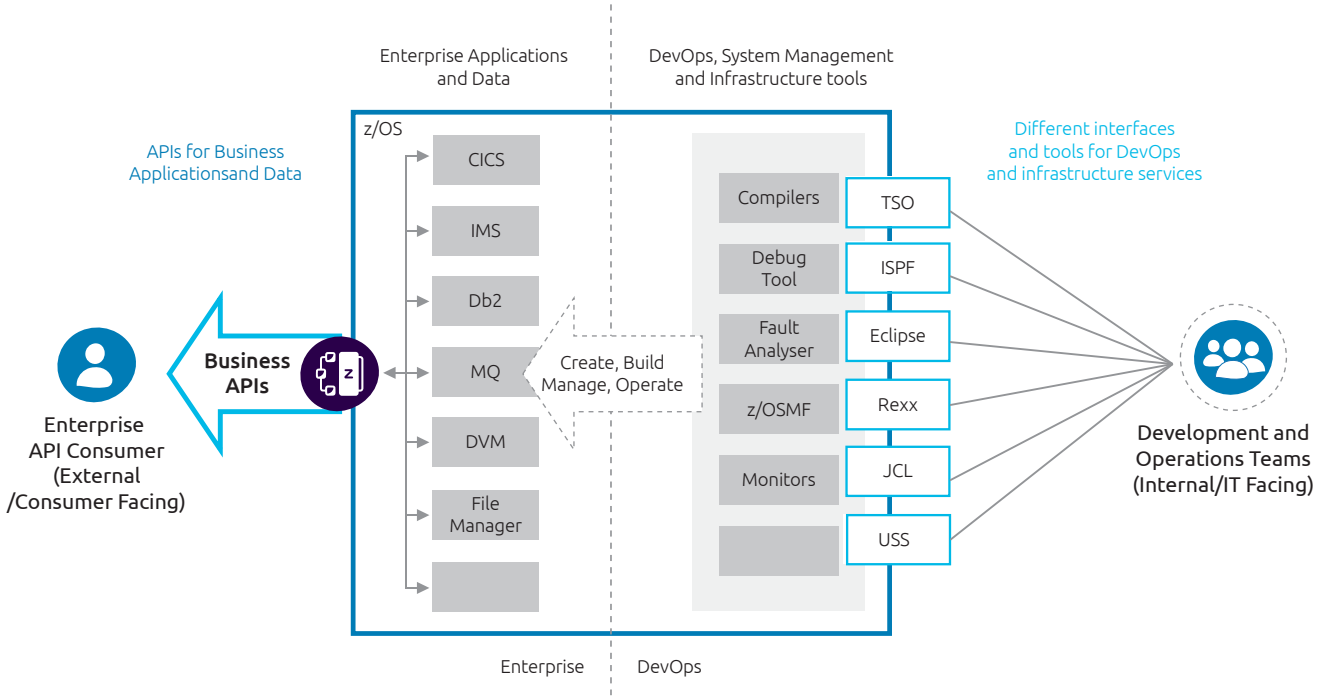

Real-time access to event driven data

Mainframe usage is specifically high for data-centric industries who want fast and secure processing of data. The vast majority of credit card transactions still run on mainframes. Unfortunately, most of this data is locked in mainframes as many organizations don't have a strategy to make this data accessible by cloud applications in real-time. Many of the organizations have implemented enterprise Datalake or Datawarehouse where the mainframe data is curated and loaded to run analytics and business reporting, but this data is mostly a day old and does not support real-time data driven decisioning.

Organizations can solve this problem by creating a Digital Integration Hub exposing core mainframe data elements in real-time to all consumers. Digital Integration Hub is an advanced application architecture that aggregates multiple back-end system of record data sources into a low-latency and scale-out, high-performance data store (Gartner definition).

IBM's z Digital Integration Hub (zDIH) provides flow of information between system of records on mainframe and cloud environment with high through-put and low latency. This is not raw data but curated information which can be accessed in real-time by cloud applications. The information can be curated from various mainframe data sources such as DB2, VSAM and IMS or directly from business applications while processing data in real-time. The data is stored in an in-memory cache which can be easily accessed using either RESTful API, Kafka topics or SQL clients. The goal of the DIH is to:

- Prevent the core applications from being overwhelmed with excessive API requests by routing majority of read-only calls to DIH cache
- Provide ability for cloud applications to consume real-time information by subscribing to Kafka topics
- Provide real-time data access through SQL client to business analysts, developers and production support team

Figure 4. z Digital Integration Hub to provide access to event driven information on mainframe

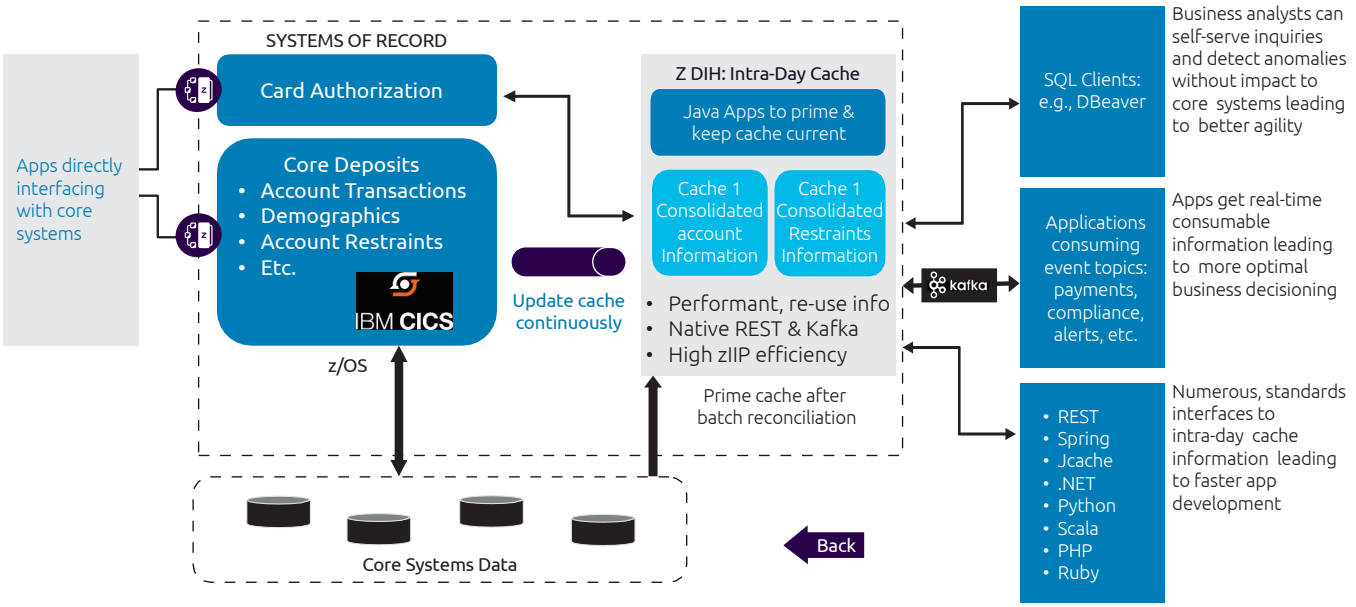

Bring AIML to data at source

Organizations are rapidly adopting usage of AIML to extract value out of a staggering amount of data being collected due to digital transformation. The AIML workloads are currently run on data curated on either on-premise or cloud hosted Datawarehouse which generally does not have real-time data.

The value extracted from real-time data will be significantly higher for some of the popular AIML processes like fraud detection, sanctions and anti-money laundering. Financial organizations are currently only able to score 10% of transactions for fraud detection in real-time due to latency involved in moving data from origination platform (mainframe) to processing platform (cloud).

After launch of z16 machines, it is possible to execute AIML models on mainframe using Watson Machine Learning for z/OS (WMLz). Also, with the launch of z16, IBM has introduced on-chip AI hardware acceleration that can handle 300 billion inference requests per day with one millisecond of latency(1).

This makes it possible for credit card companies to detect fraud while in-process.

This modernization approach focuses on executing already (externally) trained models on mainframe, the development and training of models can happen on any of the cloud platform. This could have game changing implications on how AIML models are developed, trained and executed in future.

(1)Performance result is extrapolated from IBM internal tests running local inference operations in an IBM z16 LPAR with 48 IFLs and 128 GB memory onUbuntu 20.04 (SMT mode) using a synthetic credit card fraud detection model (https://github.com/IBM/ai -on-X-fraud-detection) exploiting the Integrated Accelerator for AI. The benchmark was running with 8 parallel threads each pinned to the first core of a different chip. The Iscpu command was used to identify the core-chip topology. A batch size of 128 inference operations was used. Results were also reproduced using a z/OS V2R4 LPAR with 24 CPs and 256GB memory on IBM z16. The same credit card fraud detection model was used. The benchmark was executed with a single thread performing inference operations. A batch size of 128 inference operations was used. Results may vary.

Figure 5. Train anywhere and deploy on z workflow for AIML models

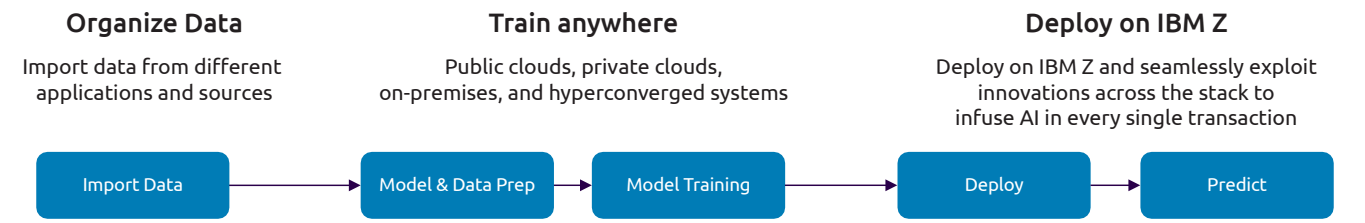

CONCLUSION

Mainframe modernization does not have to be migrating away from mainframe. Capgemini and IBM are bringing a balanced approach to help customers on their modernization journey to choose the right path for them between migration v/s in-place modernization.

AUTHOR

Nitin Chavan
is a Senior Director and Cloud Champion with Capgemini. His focus area is to help customers to adopt to hybrid cloud journey using IBM and Red Hat technologies including Z system modernization. He has 19 years of experience in providing technology solutions and has led multiple large transformation engagements for Z application and data modernization for financial services customers.

Dr. Gururaj Rao
is an IBM Fellow & GSI Ecosystem Executive at IBM Systems Group leading partnerships with worldwide Systems Integrators to expand the coverage to grow z Systems. He has been a technology and technical leader establishing the technical direction for IBM's Enterprise Systems working with inter-divisional teams of hardware, software and Research engineers and working with enterprise customers and IBM's client facing units.

About Capgemini

Capgemini is a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided everyday by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of over 340,000 team members in more than 50 countries. With its strong 55-year heritage and deep industry expertise, Capgemini is trusted by its clients to address the entire breadth of their business needs, from strategy and design to operations, fueled by the fast evolving and innovative world of cloud, data, AI, connectivity, software, digital engineering and platforms. The Group reported in 2021 global revenues of €18 billion.

Get the Future You Want | www.capgemini.com