


precisely

Best Practices for Mainframe Modernization





Mainframe systems have been the ultimate platform for powerful and secure computing for decades. They provide the core IT infrastructure upon which the world's preeminent global enterprises and largest governments depend to run their entire operations.

But in recent years, cloud computing platforms have advanced exponentially in power and flexibility, making them no less attractive to enterprise-scale businesses. Cloud platforms' elastic and infinitely scalable processing can significantly reduce IT infrastructure capital costs while enabling more efficient and responsive operations and development cycles.

Yet, for those organizations that have long relied upon the mainframe for their operations, the excitement generated by the possibility of adopting a comprehensive cloud-first IT strategy is immediately tempered by a hard reality: It is incredibly challenging and disruptive to migrate applications and data from the mainframe to the cloud.

Fortunately, you do not need to choose between mainframe and cloud computing. By focusing on a long-term strategy of coordinated and incremental modernization across both realms, you can purposefully merge their power and advantages while keeping all your strategic and tactical IT options open.

In this eBook, we will clarify the concepts of mainframe modernization, explain the drivers behind the common mainframe modernization patterns, and provide guidance on the most important methods and tools involved when putting those patterns into practice.





The Modernization Mindset

The question of when or if to move away from mainframe computing has been debated ever since cloud computing platforms, such as Amazon Web Services (AWS), Google Cloud, Microsoft Azure, and others, achieved sufficient scale and power to be credible alternatives. The arguments for and against migrating mainframe to the cloud have certainly evolved, but in the end, neither option has proven itself to be the best and only choice.

Instead, what has emerged is that organizations that currently depend upon mainframe computing are making pragmatic strategic plans for the future by utilizing both approaches. Indeed, the nature of the question has shifted away from focusing solely on mainframe migration decisions and is now commonly framed in terms of 'mainframe modernization.'

IT executives who manage mainframe systems know full well that there is nothing back-level or outdated about the platform. Consider the following: A single IBM z16 mainframe can flawlessly handle 19 billion highly secure credit card transactions daily. The platform fully supports the development and deployment of containerized applications, boasting legendary system security while achieving 99.999% uptime annually.

More directly, it is not the IBM mainframe platform that needs to be modernized but rather how and where it is integrated within your organization's increasingly cloud-first IT strategies.

Following the example of many leading global enterprises, your mission is to establish your near- and long-term business requirements and to overcome the technical challenges typically involved in building a harmonized, hybrid mainframe and cloud IT architecture. The key is knowing when and how to apply the methods and tools available for mainframe modernization and how best to sequence them across near-term projects to ensure you achieve your immediate goals without cutting off your future options.

Mainframe modernization patterns

While constantly changing business requirements, growing competitive pressures, and the emergence of powerful new technology options all certainly explain why you are making plans for mainframe modernization, exactly how you should do so is primarily guided by the structure and scale of your current cloud deployments.

Because cloud computing offers such a broad and flexible range of data storage and streaming options and virtually unlimited possibilities for running highly integrated applications and services, your particular cloud environment is undeniably unique. However, cloud deployments by large, enterprise-scale businesses tend to focus on enabling advanced and scalable data storage and access capabilities, online transaction processing (OLTP), Online BI and analytics (OLAP), and software and systems development environments.

In turn, the process of modernizing mainframe systems, that is, integrating and leveraging the mainframe's strengths and power within a cloud environment, tends to follow a few common patterns:

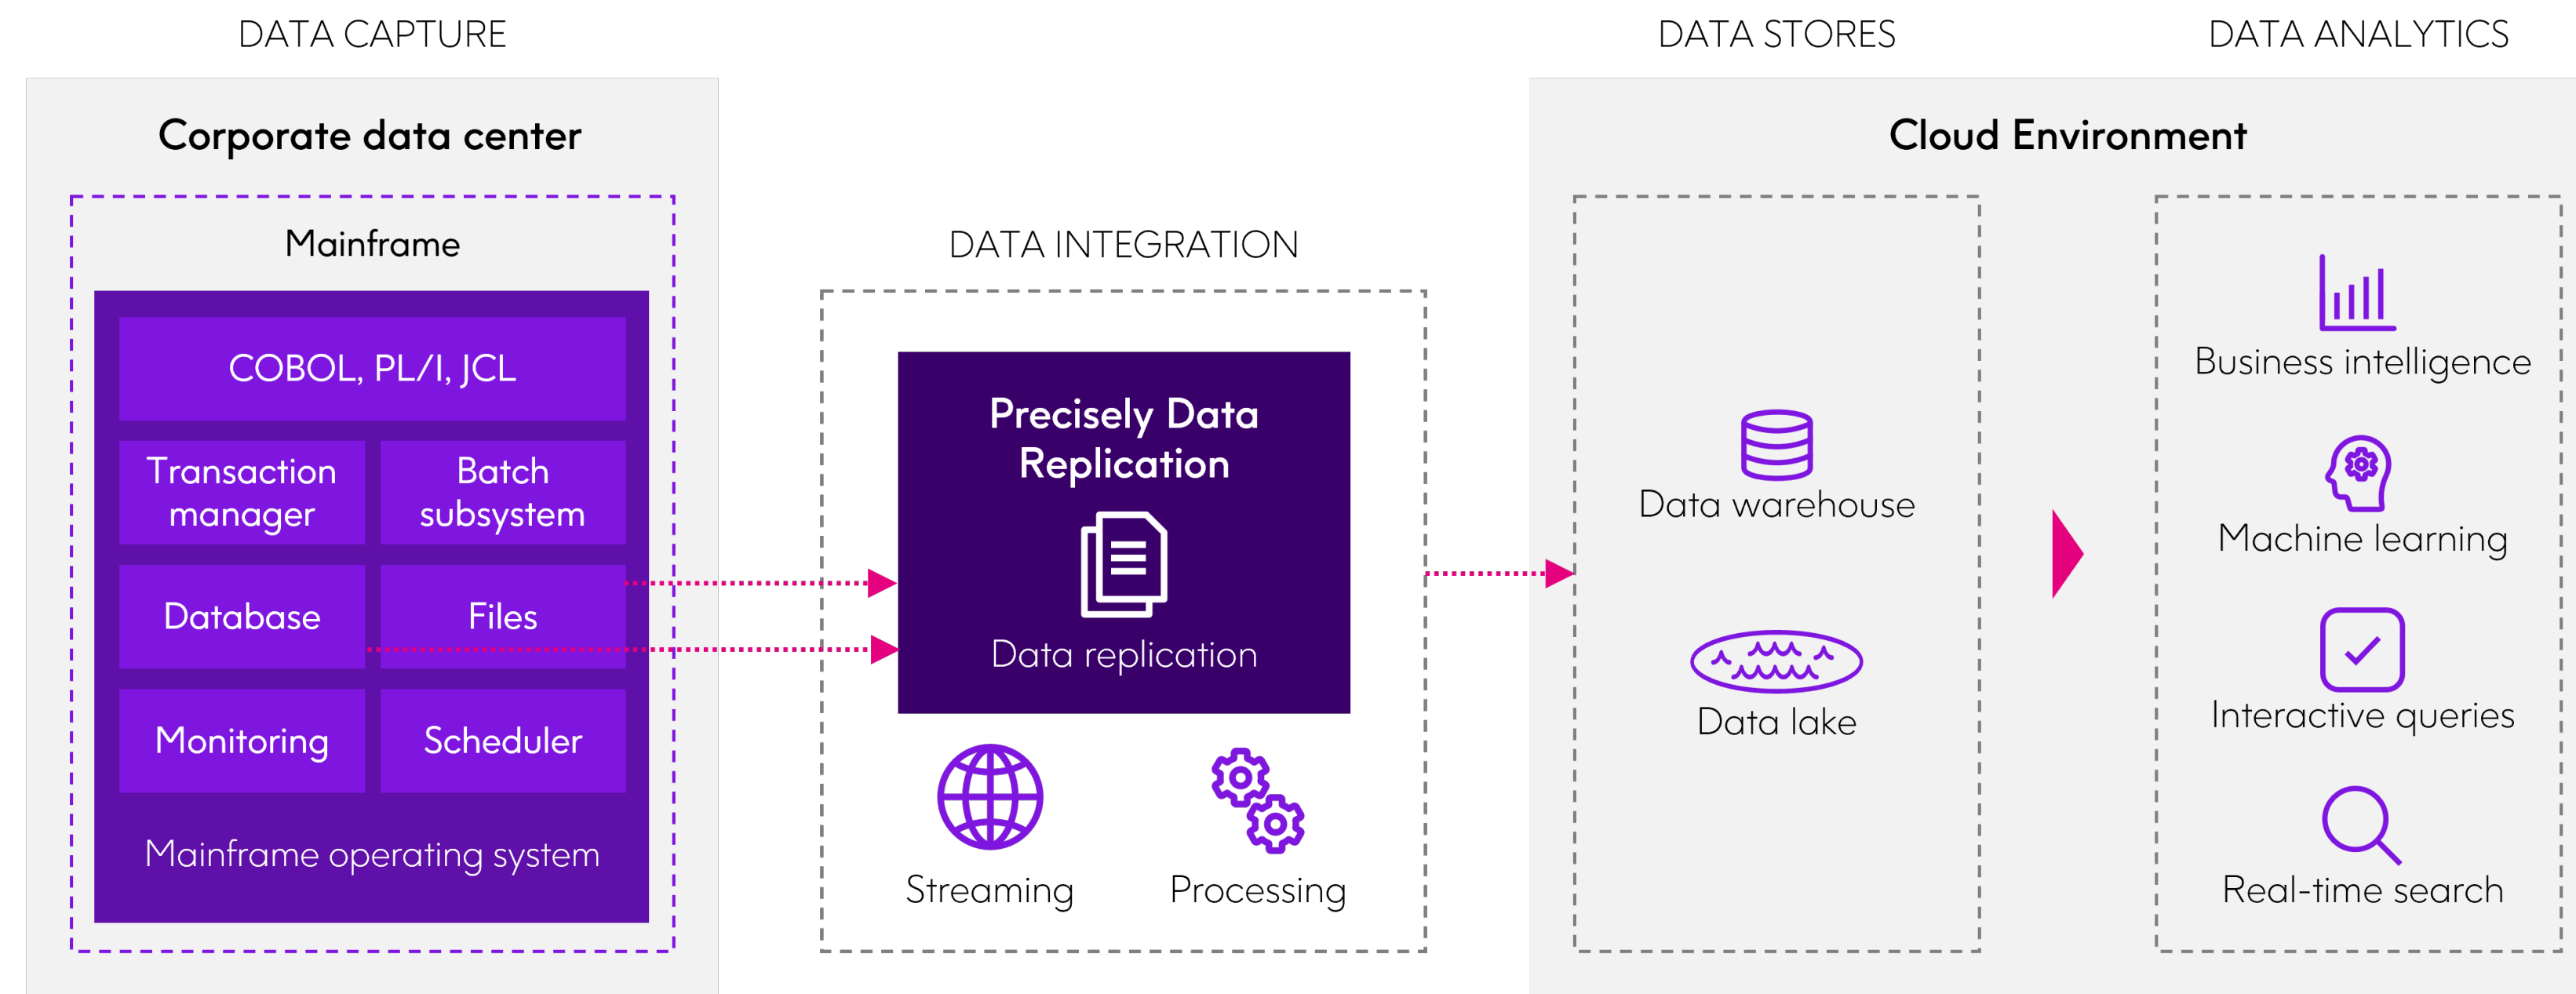
- 1. Data Replication:** Replicating mainframe data to cloud data warehouses and data lakes
- 2. Cloud-based DevOps:** Development of cloud applications and services which depend upon mainframe data
- 3. Mainframe migration:** Long-term, step-wise refactoring and replatforming of mainframe applications to cloud



The Data Replication pattern

When the most important objective is to be able to use mainframe data residing in Db2, VSAM, and IMS in the cloud, modernization is defined as accessing, transforming, and then delivering your data to cloud data warehouses and data lakes. End goals in this case include leveraging mainframe data for analytics and decision support, including business intelligence, interactive queries, and real-time search.

Mainframe data is also of exceptional value for machine learning and AI platforms. And cloud repositories certainly provide scalable, resilient, and secure data backup and long-term archive storage at a much lower cost than Direct Access Storage Device (DASD) systems.



In each case, data generally flows in one direction, replicating from the mainframe to the cloud. And although frequency and speed of updates to cloud data stores may not be as critical for such applications as they are for online transaction processing, real-time or near real-time delivery is still the standard.

To achieve the lowest possible replication latencies, avoid replicating directly to your cloud database systems. The reason is that each transaction to a database requires many I/O instructions and acknowledgments to complete. And the more distant your cloud data center is from your mainframe, the greater the impact of each transaction on replication latency. These two factors quickly add up to unacceptably long latencies for any purpose.

Instead, replication to event streaming platforms, such as Kafka, Amazon Kinesis, Rabbit MQ, etc., will achieve far lower latencies. Event streaming platforms enable a continuous “firehose” approach to replication, with throughput rates as high as tens of thousands of records per second, compared with a few thousand when using direct-to-database replication.

Canadian Life Insurance Company

Objective:

Giving insurance providers and members fast and reliable online access to information about their claims

Challenges:

- Delivering claims data stored in multiple mainframe Db2 and IMS repositories to cloud-hosted systems
- Current batch/ETL delivery process cumbersome and far too slow
- Migration or refactoring applications totally out of scope for cost and deployment timeframes

Solution:

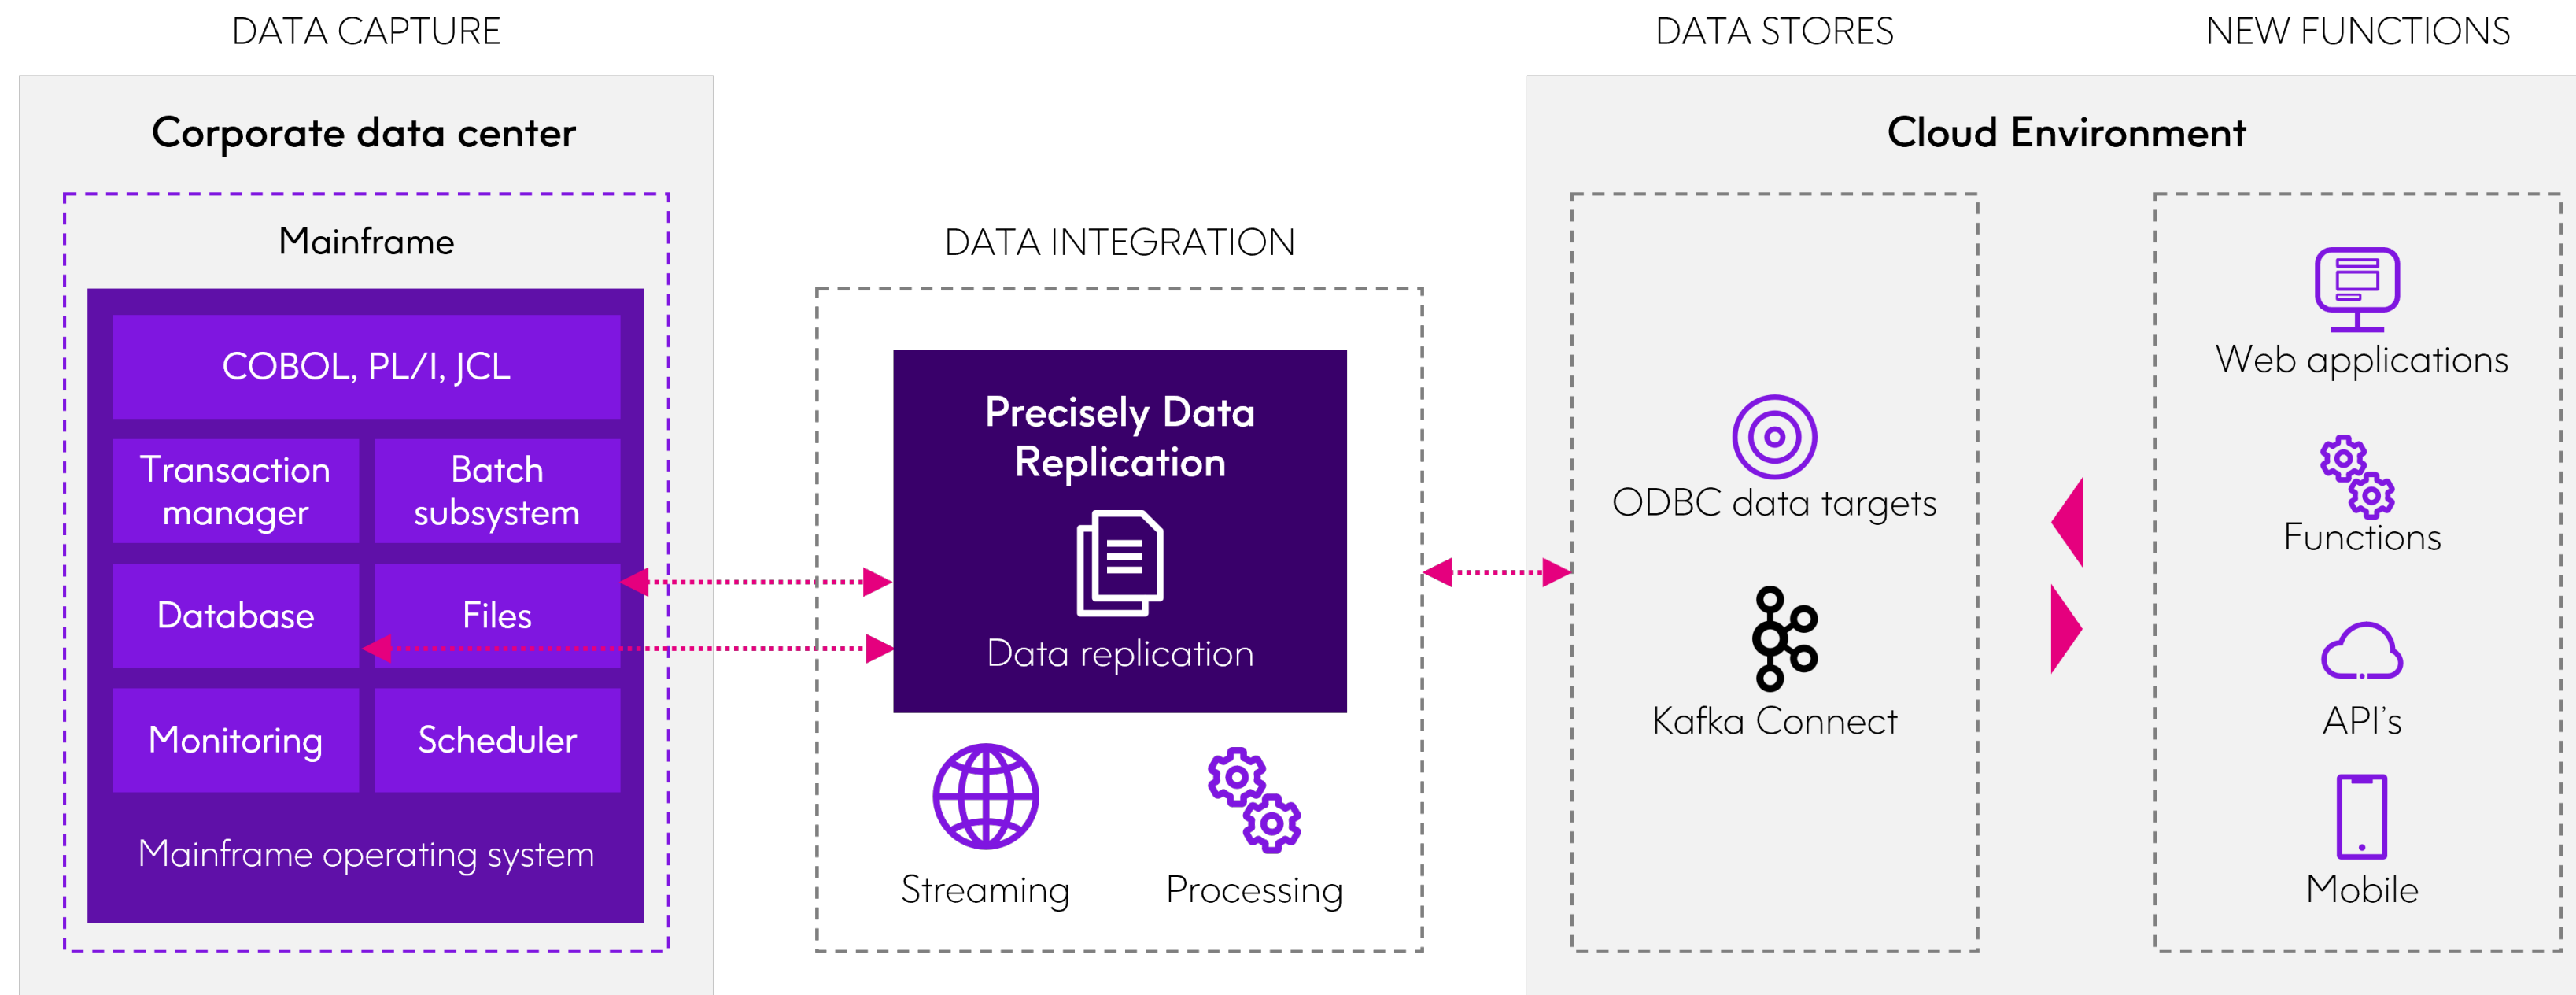
Precisely Connect for real time delivery of mainframe-hosted Db2 and IMS data to Kafka for ingest and processing in AWS-hosted cloud applications

Results:

- Mainframe Db2 and IMS data fully integrated into cloud-hosted web services without disruptive and costly migration
- Far lower MIPS and system costs, minimized programming and administrative burden than prior Batch/ETL processing solution

The Cloud-based DevOps pattern

A more advanced and, by default, more involved mainframe modernization objective is integrating both mainframe data and mainframe development efforts with cloud-based DevOps. In this case, the end game is to build and deploy more powerful and tightly aligned web applications, system APIs and functions, mobile apps and UIs, etc. Ultimately, these goals are achieved by doing everything possible to eliminate delays and disconnects between your increasingly interdependent cloud and mainframe development efforts.



In the near term, blending mainframe development team members and their processes into all cloud application and service development efforts will require substantial and consistent executive support, professional skills upgrades, and change management disciplines. But over time, as mainframe and cloud development efforts are more fully synchronized and strengthened by having hybrid teams, deployment costs and timeframes will be reduced while code quality and system interoperability will naturally improve.

For any organization whose longer-term mainframe strategy includes, or even potentially includes, actively refactoring or fully replatforming mainframe applications, starting the process of unifying and strengthening DevOps now will ensure a stronger, more migration-ready stance that will simplify and accelerate migration planning and execution.

North American Bank

Objective:

Delivering online 'digital banking' to millions of customer

Challenges:

- Essential data hosted in extensive mainframe VSAM repositories
- Delivery of mainframe data for use across multiple online channels and applications
- Build and deployment of new cloud channels and services hobbled by mainframe access issues.

Solution:

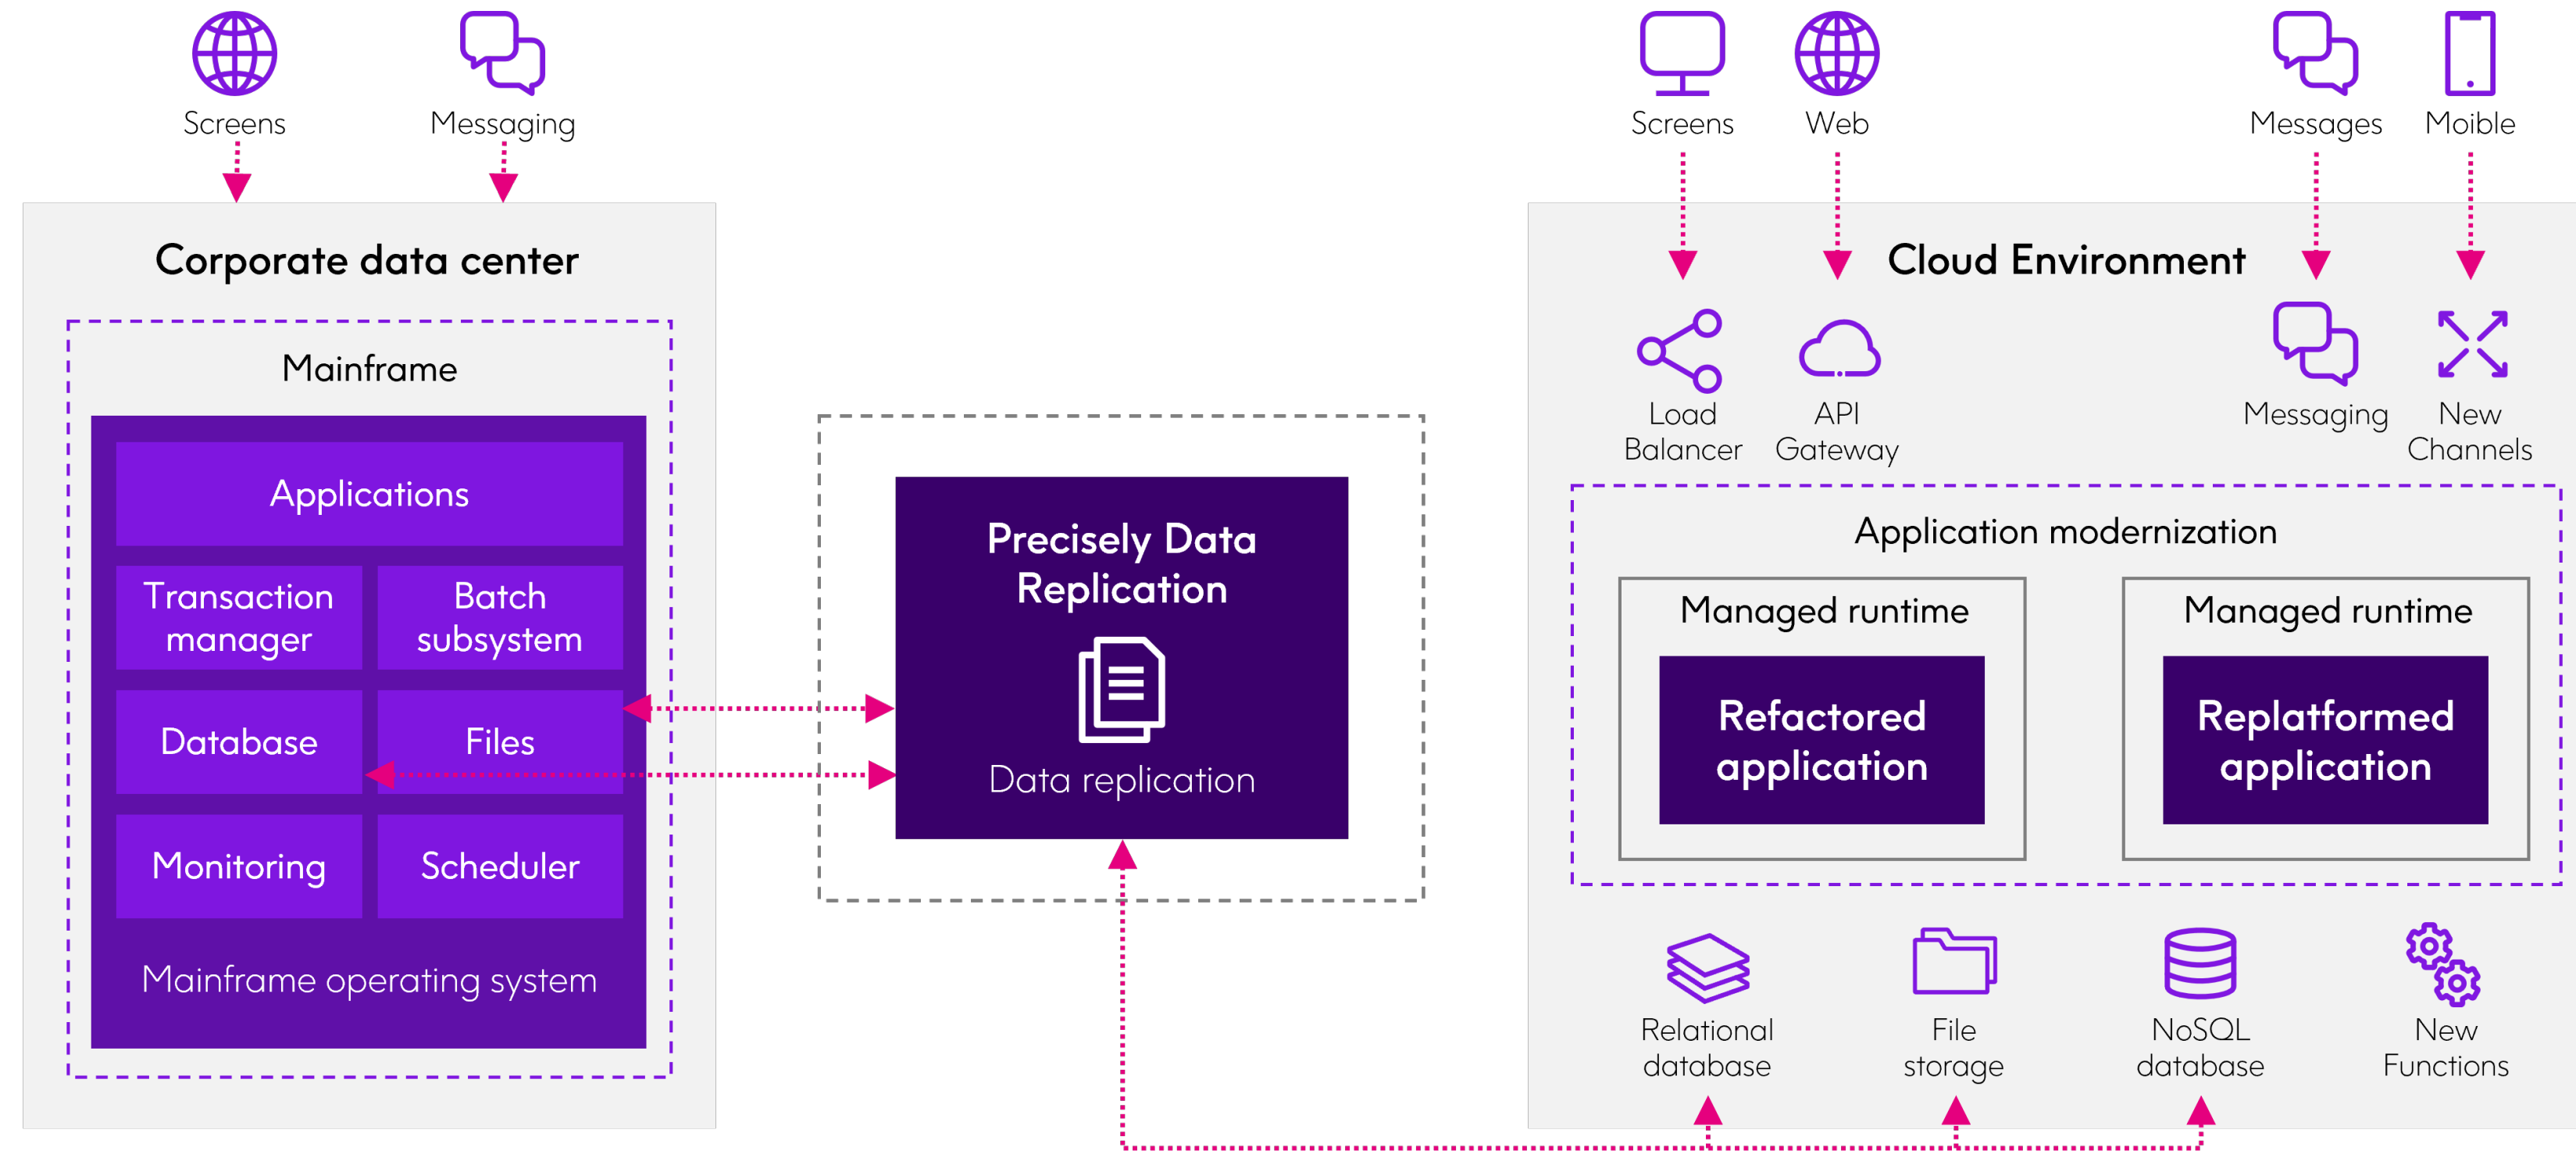
Precisely Connect for real time delivery of VSAM data to Kafka for highly distributed processing across AWS-hosted applications.

Results:

- Vastly improved online experience driving gains in customer recruitment and loyalty
- Streamlined development cycles, faster time-to-market than competitive alternatives (150-250 days)
- Elimination of inefficient and redundant mainframe system processing and data administration burdens.

The Mainframe Migration pattern

When the commitment is to migration, that is, to fully replace most or all mainframe applications and processing with cloud-native systems, a more comprehensive approach is needed. Mainframe development efforts and operations management must be fully coordinated and carefully managed. This is because migration from the mainframe will involve running your business on a constantly evolving IT base, maintaining uncompromised mainframe processing while simultaneously building a parallel environment in the cloud, one application and process at a time.



Then, as individual parts of your currently trusted, proven mainframe processing are handed off to their cloud successors, every interdependency they have with the remaining mainframe processing functions, including zOS and its deeply embedded security functions, system logs, metadata, and even such arcane structures as Copybooks, must be faithfully maintained.

This interconnection and interdependency must be engineered and maintained over many months, possibly years. Even the largest organizations with a full bench of highly skilled and cross-trained professionals will inevitably need to rely heavily on professional services and support from their cloud provider and mainframe platform experts to see the mission through to its successful conclusion. Cloud platform-native solutions and services, such as the Mainframe Modernization Service offered by AWS, will be invaluable for any organization taking on large-scale, migration-focused mainframe modernization projects.

North American Payments Processor

Objective:

Moving clients' core, mission-critical payment and transfers applications to AWS cloud platform

Assured scalability and easier application maintenance

Challenges:

- Real-time data access without impacting mainframe operations
- Repeatable yet flexibly customizable deployment processes for data integration

Solution:

Precisely Connect for real-time delivery of IMS, VSAM and Db2 data to multiple clients' AWS environments via Kafka

Results:

- Fully scalable, near-real-time data delivery
- No backpressure on core mainframe-hosted banking system processing
- Utility-like implementation and operation speeds and simplifies DevOps for rapid client onboarding

Essential capabilities for mainframe modernization

No matter what business imperatives or technological advancements are driving your mainframe modernization decisions, and regardless of how quickly or thoroughly you plan to move your current mainframe capabilities into the cloud, some foundational IT capabilities must be in place to get the job done, and several critical IT operations and business policy issues must be addressed.

To help ensure that your specific project plans are properly structured and executed to support strategic success, here are the top-line requirements and proven best practices for mainframe modernization, drawn from Precisely's decades of leadership in mainframe systems technology and the vast technical and implementation experience of major cloud providers such as AWS.



First and foremost: mainframe data replication

To modernize the mainframe, large volumes of data must be extracted from System z platform storage devices and delivered to cloud storage systems and data streams. While that requirement may seem clear and straightforward, fulfilling it is complex. So, to begin with, here are some key technical and tactical factors for mainframe data replication.

Data extraction and transformation

For data stored in DASD, extraction entails much more than a direct read/write operation. Extended Binary Coded Decimal Interchange (EBCDIC), packed decimal, and zoned decimal data will require unpacking and translation to Unicode and other formats to be useable in your cloud-based systems. In addition, at the record level, mainframe data is stored in mainframe-specific fixed, VSAM, Fujitsu, or Mainframe Variable flat file formats, all of which must be transformed before delivery to cloud data storage systems.

While cloud-hosted relational database tables are strictly structured in rows and columns, nothing enforces a set data structure on the mainframe. Instead, sections of COBOL code within programs, called Copybooks, define how the data should appear. So, to be useable in cloud systems, data records sourced from the mainframe must be changed in several ways, with the inevitable result that they no longer match up properly with the original Copybook-defined records. This mismatch causes problems, including seriously interfering with the bi-directional data replication between cloud and mainframe required under a subset of mainframe modernization patterns.

So reliable and efficient methods for reencoding and mapping data across mainframe and cloud environments must be in place.

Life Insurance Company

Objective:

Global-scale digital transformation to apply advanced analytics and machine learning (ML) to call-center operations

Challenges:

- Delivering huge volumes of disparate data from mainframe and streamed sources to cloud analytics and ML systems
- High costs and time requirements for developing and deploying custom APIs for mainframe

Solution:

Precisely Connect for delivery of mainframe, streamed audio, and multiple other data sources to AWS cloud-hosted Confluent Kafka platform.

Results:

- Elimination of major mainframe API development and maintenance workloads
- Secure delivery of streamed call center audio files to cloud analytics
- Integrated processing of mainframe and other data in AWS cloud analytics and ML



Changed data capture

While efficient batch data transfer operations will still be required as part of any mainframe modernization project, the most significant data movement is the real-time replication of changes made to data on the source systems to the target environment. Due to all the data format and storage structure differences just discussed and numerous other mainframe OS and hardware factors, the processes for recognizing when data has changed are quite complex. And on a platform where a single server can process billions of transactions daily, all the processes involved in identifying and capturing data changes and subsequently transforming and transmitting them need to be as fast.

For any organization, no matter how large or capable their mainframe application development teams may be, building bespoke, in-house mainframe data replication solutions makes little sense. The only practical and cost-efficient option is to adopt tools and applications from leading mainframe replication solution experts.

Metadata / Data provenance

When data is moved, accessed, or replicated between systems, metadata regarding data lineage and changes must be recorded and replicated. This chronological record becomes critical for recovery if data is corrupted, lost, duplicated, or rendered incorrect or unusable. It is also essential for regulatory compliance, including annual financial auditing, data privacy controls, and myriad governmental laws and regulations.

When moving mainframe data into your cloud computing environment and from cloud to mainframe when bi-directional replication is required, metadata management is critical. Data will move frequently between multiple systems and applications within your cloud environment, and the initial processes involved in making mainframe data available to your cloud environment will create a huge new volume of metadata records that must be fully safeguarded.

This critical requirement argues against the in-house development of bespoke mainframe replication solutions.

Secure data replication

Within the mainframe, data security is managed by IBM RACF, Broadcom ACF2, or Broadcom TOP SECRET products. Mainframe modernization requires your SEIM solution(s) to integrate with mainframe security processes. Continuous and fully integrated access to security records is essential for your operations management and regulatory compliance teams.

Though it may seem completely obvious, it still bears mentioning that industry-leading security protocols and solutions are applied to the staging, transmission, and ingestion of all data moving between mainframe and cloud. Again, the solutions and expertise offered by your cloud environment provider and mainframe technology partners provide invaluable protection. The best solutions will be those that your cloud and mainframe partners mutually and fully support.

Scalable, repeatable replication schemas

As discussed earlier, your mainframe modernization strategy will be executed through an ongoing series of projects, which will involve managing data replication from your mainframe systems to your ever-evolving cloud environment and potentially from cloud data sources back to the mainframe. Creating and frequently modifying the many data replication models crucial for directing all your data flows amid continuous changes to both platforms will be an unavoidable, ongoing workload for your development team.

But over time, the constant cycles of replication schema rework can become a tedious and unwelcome chore that interferes with and delays the more critical (and professionally interesting) development work your team needs to complete. And, as is always the case, time is money.

This aspect of mainframe modernization projects is frequently overlooked. However, the issue can be avoided entirely by utilizing a replication solution that simplifies and automates the building, modification, and re-use of its underlying replication models.



Seeing the whole chessboard

Beyond considering the technical details of mainframe data integration covered above, a discussion of higher-order IT management issues is also warranted.

Like all large, enterprise-scale operations, your organization is undoubtedly well along in its efforts to achieve maturity across several key IT and general business management disciplines. You may still be working to improve some aspects of Data Governance and Data Quality, end-to-end enterprise data security, HA/DR resilience, etc.

For other newer mandates, such as ESG tracking and reporting, adoption of advanced AI technologies, Geo Addressing and Data Enrichment, Master Data Management, etc., you may be at an earlier stage of your journey, even possibly just beginning to implement the required solutions and business processes.

As impactful as all these efforts are on your organization, mainframe modernization is in a class by itself. Expanding and remodeling the house you live in is one thing. But making significant changes to its foundation, and even picking it up and moving it to a new plot of land? That is a much bigger event that directly impacts all your other projects and responsibilities. The chess master's oft-repeated admonishment to "See the whole chess board" is appropriate in this case.

Data Quality

Maintaining data quality is a fundamental part of enterprise IT management. You invest heavily in time, staffing, and tools to guard against the many ways data can become incorrect, outdated, out of synch across data stores, and outright corrupted and useless. Mainframe modernization unavoidably puts your data at increased risk because every time data is accessed, manipulated, reencoded, transmitted, or distributed across new and additional storage systems, applications, and users, the opportunities for damage to that data multiply quickly.

So, as you execute your mainframe modernization plans, you and your entire team need to be aware of and actively mitigate the added potential for such pervasive change to inject chaos into your data.

Data Governance

A similar situation exists with regard to your data governance practices. While you are involved in moving and even re-structuring great swathes of your enterprise data, keeping track of exactly where it is stored, who has access to it, where and how it is being utilized... the effort involved in driving every data governance policy and process you have in place will naturally increase. Simultaneously, many new policies and processes will need to be developed and implemented to extend data governance to new data stores, systems, and users mainframe modernization is creating.

Data Observability

If your organization has not yet invested in solutions and tools for comprehensive data observability, and even if you have done so, know this:

At all times, but especially when applying wholesale changes to your data stores, data flows, and applications, comprehensive data observability is mission-critical. Only through highly automated, deep, and continuous monitoring and alerting will you be able to identify and manage the potentially vast increase in data and operations issues that your mainframe modernization project will create.

Address data consistency and enhancement

Accurate and complete address data is essential for customer-facing applications and a wide range of analytics, machine learning, and AI systems. Real-time integration of high volumes of mainframe data into cloud systems and data stores will make achieving and maintaining consistency of address data across all systems even more difficult.

In addition to that fundamental concern, extensive address and geographic data enhancement processes are commonly required to support advanced analytics and AI systems. Especially when multiple such applications and services are involved, it makes sense to centralize and unify enhancement processing upstream from the systems being served. This points out yet another reason to replicate mainframe data via streaming, as it is easier to set up subsequent data enhancement systems and services as a consumer of an event streaming platform and avoid laying more I/O burdens on your DBMS.

Undertaking any mainframe modernization project should prompt a thorough review of the methods and tools you currently use to address data preparation, enhancement, and maintenance to ensure the additional demands can be efficiently met.



Begin Here...

At this point, it should be clear that mainframe modernization presents enormous potential gains and advantages for any organization with IBM Power z systems in its IT portfolio. But it also presents a whole host of IT issues to be dealt with and potentially overwhelming operational challenges. At best, this eBook can only be a starting point for evaluating mainframe modernization options and investment requirements.

Just by itself, understanding that your mainframe modernization options go far beyond simply migrating away from the hardware can be transformative to your strategic planning. Understanding the modular, configurable nature of the tools and tactics for moving mainframe modernization forward and considering the models and patterns followed by leading global enterprises can be a helpful springboard for building your modernization scenarios.

Our review of some of the most important technical and operational challenges involved in modernization is offered in a spirit of transparency to help you develop realistic strategies and anticipate appropriate time frames and investment levels for your project plans.



Precisely powers mainframe modernization with trusted data

[Precisely Connect](#) accelerates mainframe modernization and migration projects by enabling powerful and efficient real-time changed data capture (CDC) and mainframe data replication to cloud platforms.

Precisely Connect is part of the [Precisely Data Integrity Suite](#), which runs natively in cloud-native platforms and empowers fast, confident decisions that fuel your data-driven initiatives. The Precisely Data Integrity Suite is a set of seven interoperable services that enable your business to build trust in its data. Data with integrity has maximum accuracy, consistency, and context – empowering fast, confident decisions that help you add, grow, and retain customers, move quickly and reduce costs, and manage risk and compliance.

Whether your organization's focus is improving the customer experience, automating operations, mitigating risk, or accelerating growth and profitability, The modular, interoperable Precisely Data Integrity Suite contains everything you need to deliver accurate, consistent, contextual data to your business – wherever and whenever it's needed.

[Learn more about data integration](#)



Precisely is the global leader in data integrity, providing accuracy and consistency in data for 12,000 customers in more than 100 countries, including 99 of the Fortune 100. Precisely's data integration, data quality, data governance, location intelligence, and data enrichment products power better business decisions to create better outcomes. Learn more at www.precisely.com.

www.precisely.com

©2023 Precisely. All other company and product names used herein may be the trademarks of their respective companies.