IBM Z Data Access

Central Canada Db2 Users Group

Richard Ruppel
Brand Technical Specialist
Data and AI for IBM Z
IBM Technology - Americas
rmruppel@us.ibm.com

11 – 13 May 2026

IBM

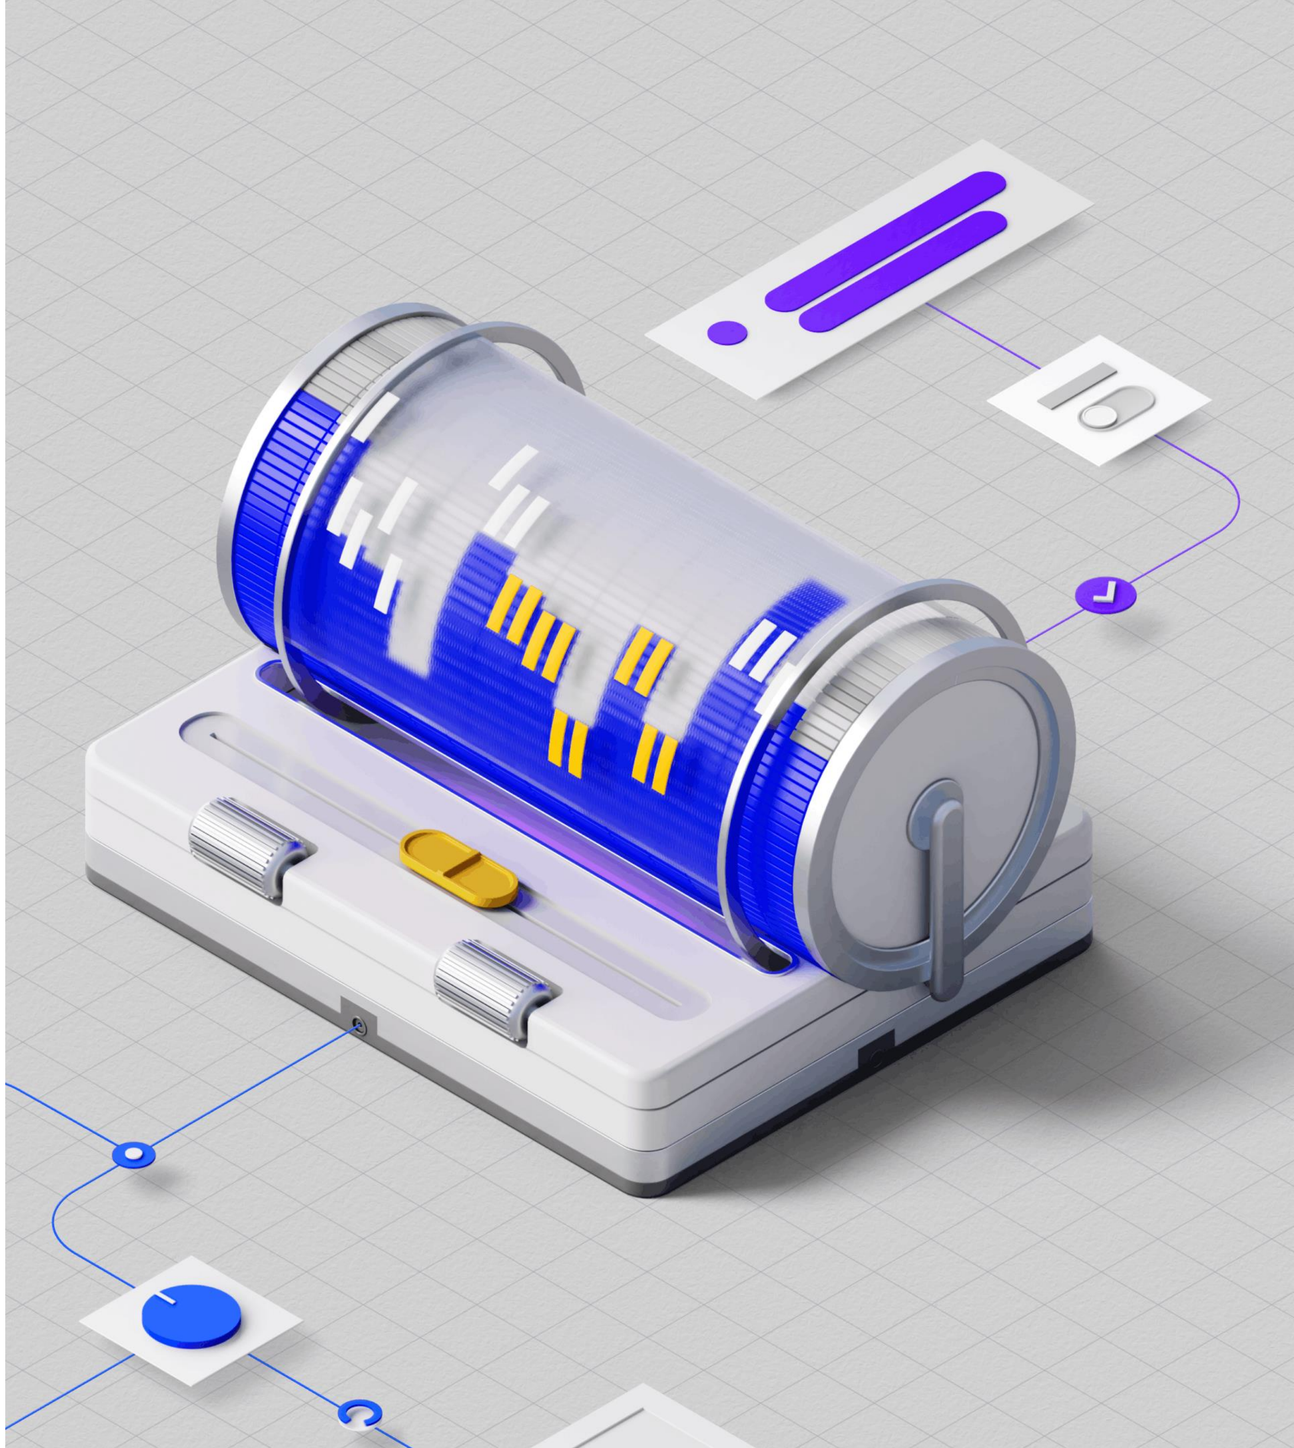

Before we talk about Data, can we talk IBM Z Technology?

- Performance
- Capacity
- Data privacy & protection
- Resiliency
- Dedicated workload accelerators
- Sustainability
- Stack optimization
- Simplification

2019 → 2025

IBM z15

5.2 GHz
5nm

Accelerated Compression

Accelerated Sort

Secure Execution

System Recovery Boost

IBM z16

5.2 GHz
7nm

Accelerated AI

Quantum-Safe System

Secure Boot

Memory Encryption

Flexible Capacity for Cyber Resiliency

IBM z17

5.5 GHz
5nm

IBM Spyre Accelerator

Multi-model, agentic, & gen AI

Accelerated I/O

Standardized Quantum-Safe Algorithms

AI-powered security

IBM z17

- 5nm Silicon Lithography*
- 8 Cores per PU Chip design
- +1 I/O Engine per chip - DPU
- Dual Chip Module packaging
- 4 PU DCMs per Drawer, up to four CPC drawers
- Integrated I/O with PCIe Gen5
- 2nd Gen IBM Accelerator for Artificial Intelligence (AIU) – on chip

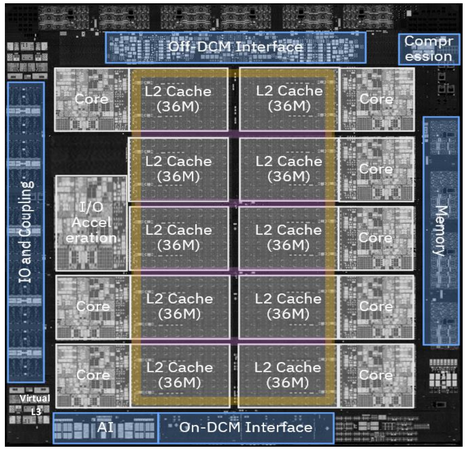

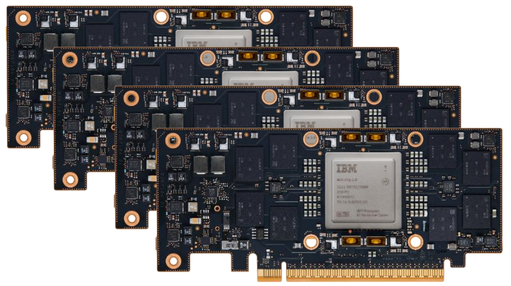

IBM Spyre Accelerator cards

- 32 additional AI cores on extended adapters, built for Gen AI models
- Adapters fit into standard PCIe slots
- Up to 48 adapters per system
- Sharding to support larger Gen AI models

What does your Data landscape look like today?

How long ago did you implement your data strategy?

- 1990s – Client / Server, start of Data Warehousing
- 2000s – Business Intelligence, web-based apps
- 2010s – Mobile access, data lakes, cloud computing
- 2020s – Analytics and AI-everything

Where did we end up?

- ➢ Complexity
- ➢ High Latency
- ➢ Security / Governance Concerns
- ➢ Increase Costs
- ➢ Did I mention Complexity?

😩 Is there anyone left that understands this?

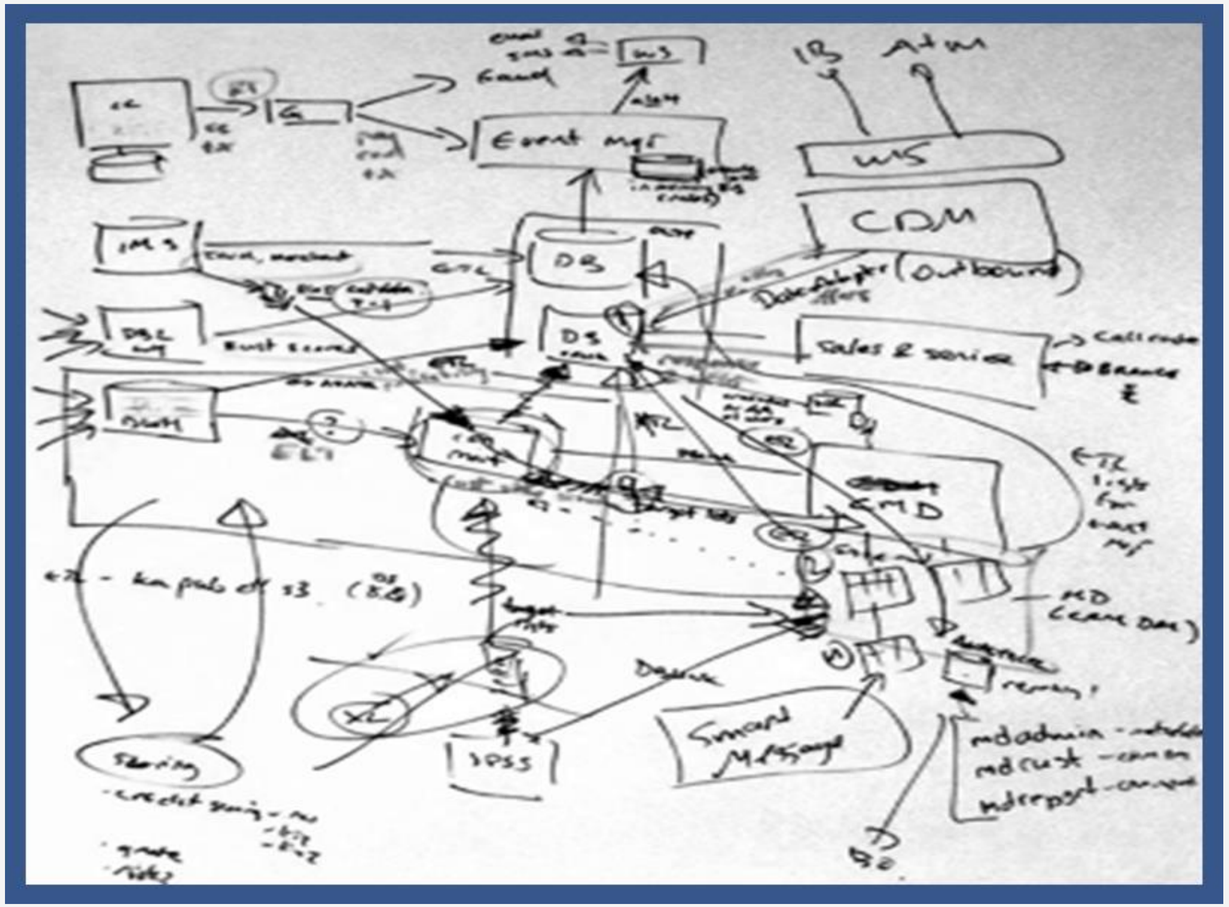

Who remembers this?

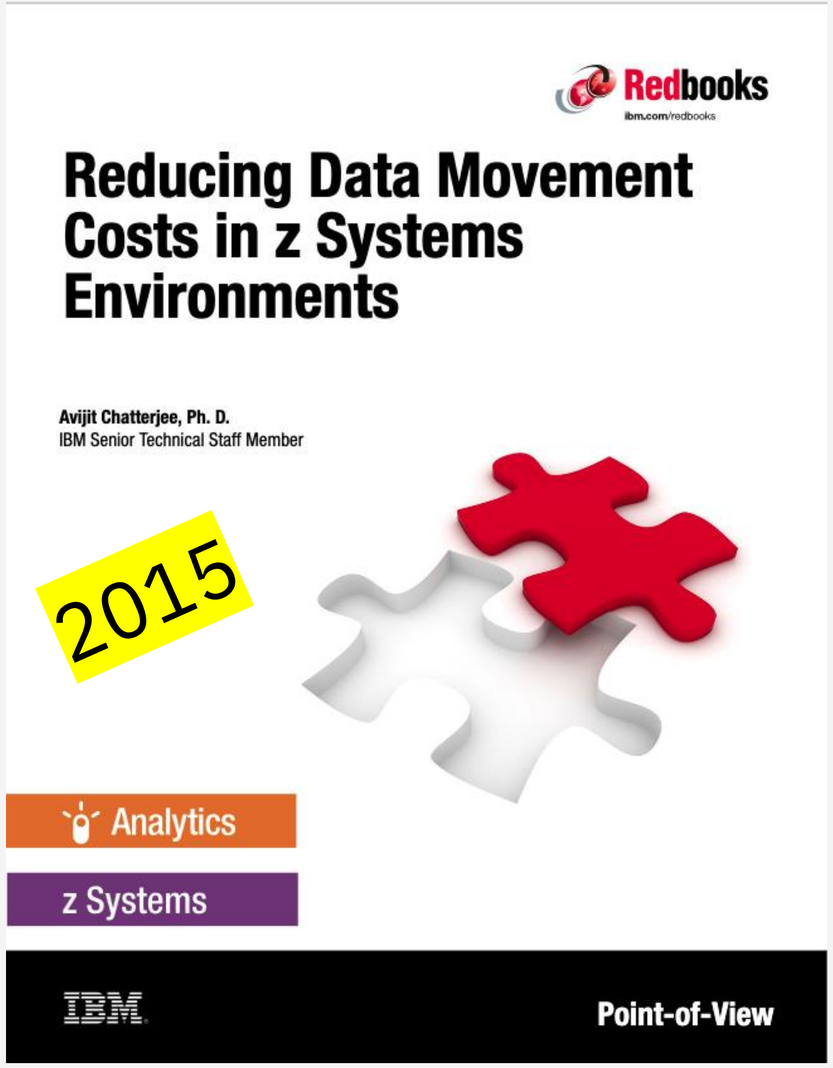

https://www.redbooks.ibm.com/redpapers/pdfs/redp5217.pdf

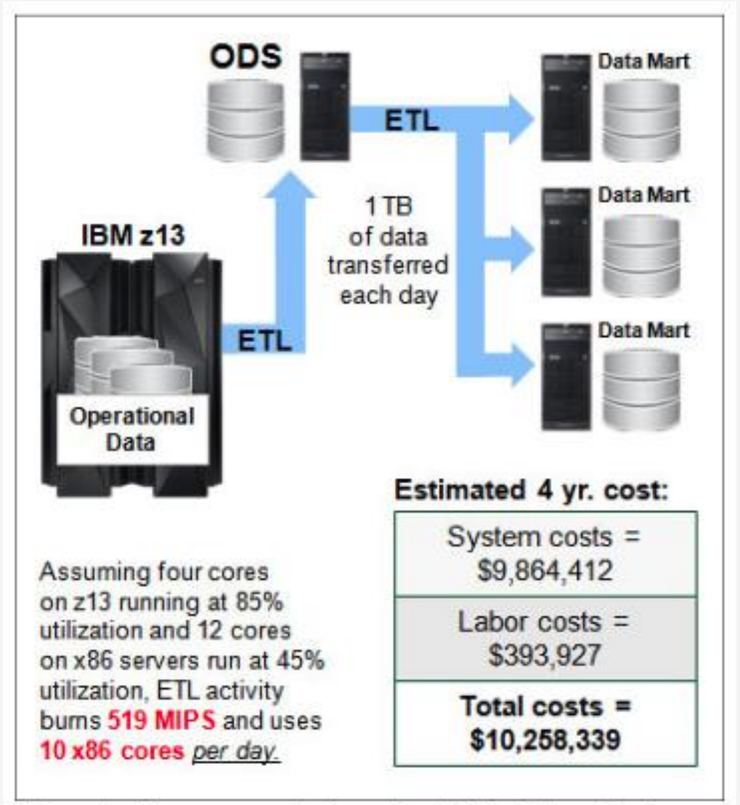

Figure 1 Four-year cost of moving 1 TB of data daily to an ODS and three data marts

The **Data Gravity** story:
Bring your applications and analytics to the data.

And now **AI Gravity** too!

Clients report that, on average, they use approximately

25% of CPU

copying their data off the mainframe...
resulting in less timely and less insightful access, analytics, reporting and warehousing.

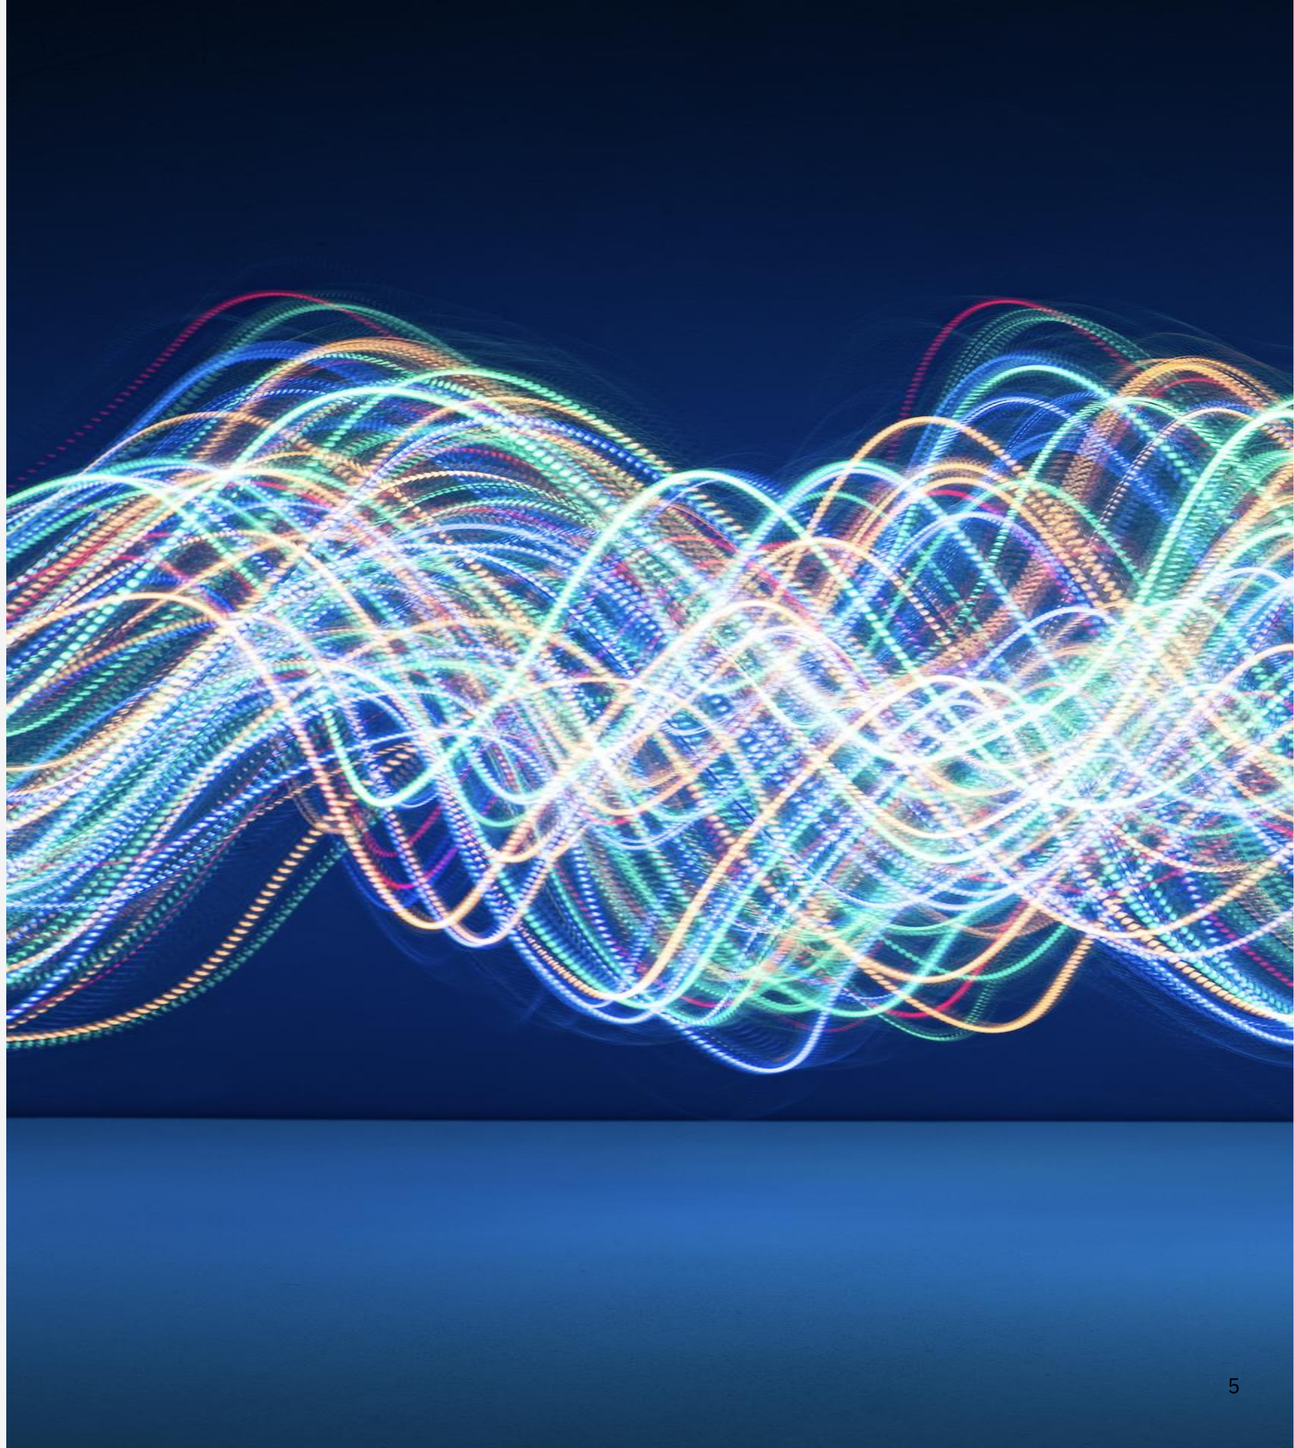

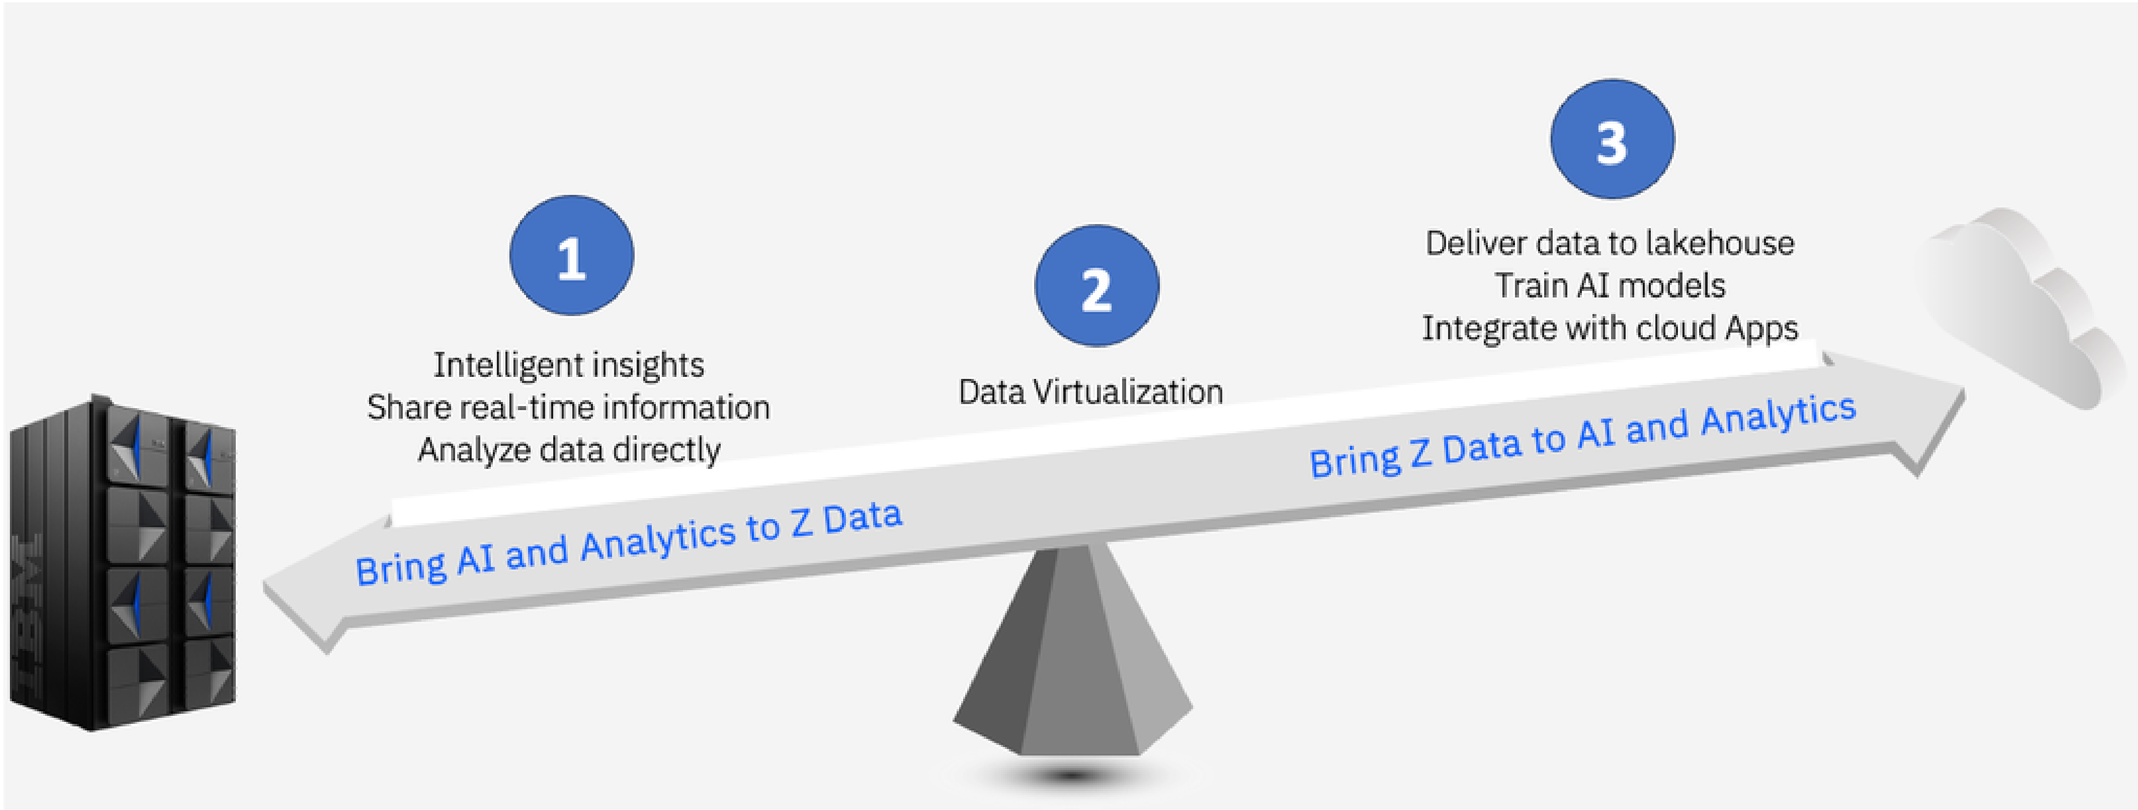

1. Bring Analytics & AI to the Data on the mainframe
   When freshness, security, and low latency matter, analytics run directly on Z.
2. Keep Data on the mainframe and Expose It Virtually
   Virtualization delivers current data to modern applications without replicating the system of record.
3. Selectively Deliver Z Data to Cloud for Broad AI Use
   For large-scale analytics and model training, IBM enables governed replication.

Hybrid Analytical and Transactional Processing - Db2 for z/OS and IDAA

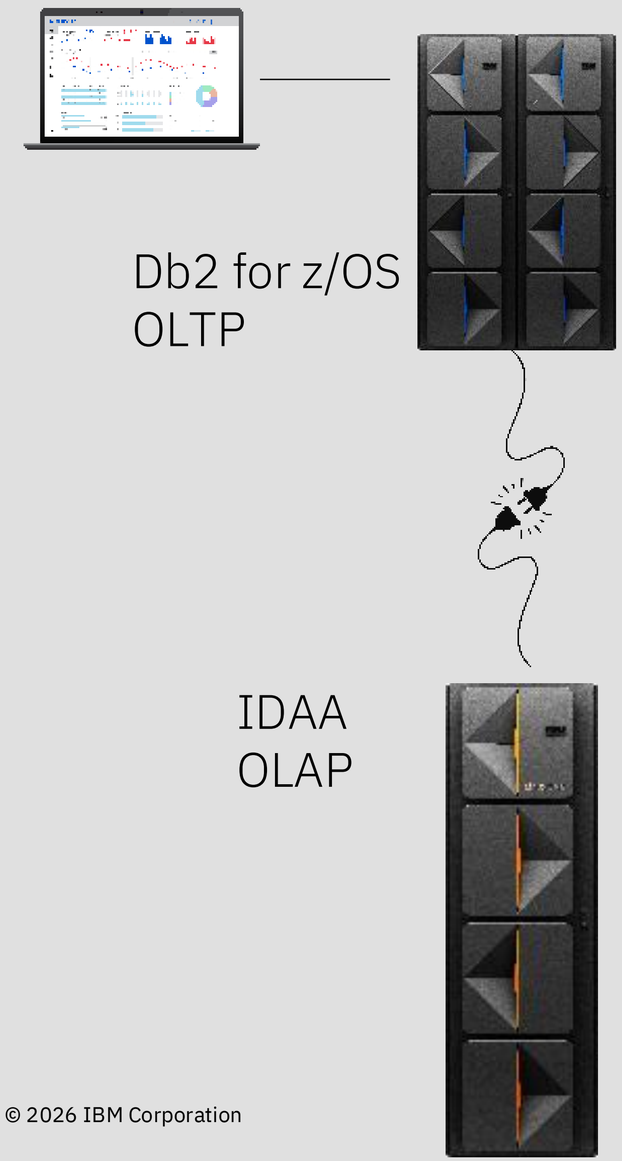

Ultimate performance

Db2 for z/OS as row-based data store for transactional access

IDAA as columnar-based data store for analytical access

Complete workload isolation

no impact to transactional workload

choice of scale-up or scale-out

Unlimited scaling

unique in industry: scale-up and scale-out

does not require application awareness

New Version – V8

LOB Support

Accel to Accel Copy

Use cases

Massive query acceleration

Online Archiving

ETL/ELT and in-database analytics acceleration

IDAA only data with accelerator only tables

Virtual data integration: fast federated joins across Db2 for z/OS systems

Load IMS/VSAM/SMF data into IDAA via IDAA Loader

No migration for existing workloads

Industry-unique, patented data coherency protocol

Integrated Synchronization

SQLDI Pro and seamless Integration with IDAA

SQLDI Pro works with IDAA for cost- efficient training and inference

Bringing AI to Where Enterprise Z Data Lives
SQL Data Insights – Built into Db2 13 for z/OS

Ease of Use

- Build Neural Network powered relationship maps using unsupervised training over (unlabeled) structured data
- Simply select data, enable training and Db2 for z/OS builds a data relationship model
- Apply relationship maps and built-in AI-related functions within any SQL statement
- Readily interpret underlying reasons for insight

Major Benefits

- No deep data scientist skills required
- Rapid time to develop and deploy AI
- No specialized architecture
- Efficient AI scoring (elapsed time, CPU, throughput)
- Highly efficient retraining and redeployment
- No data latency
- Model can address multiple questions

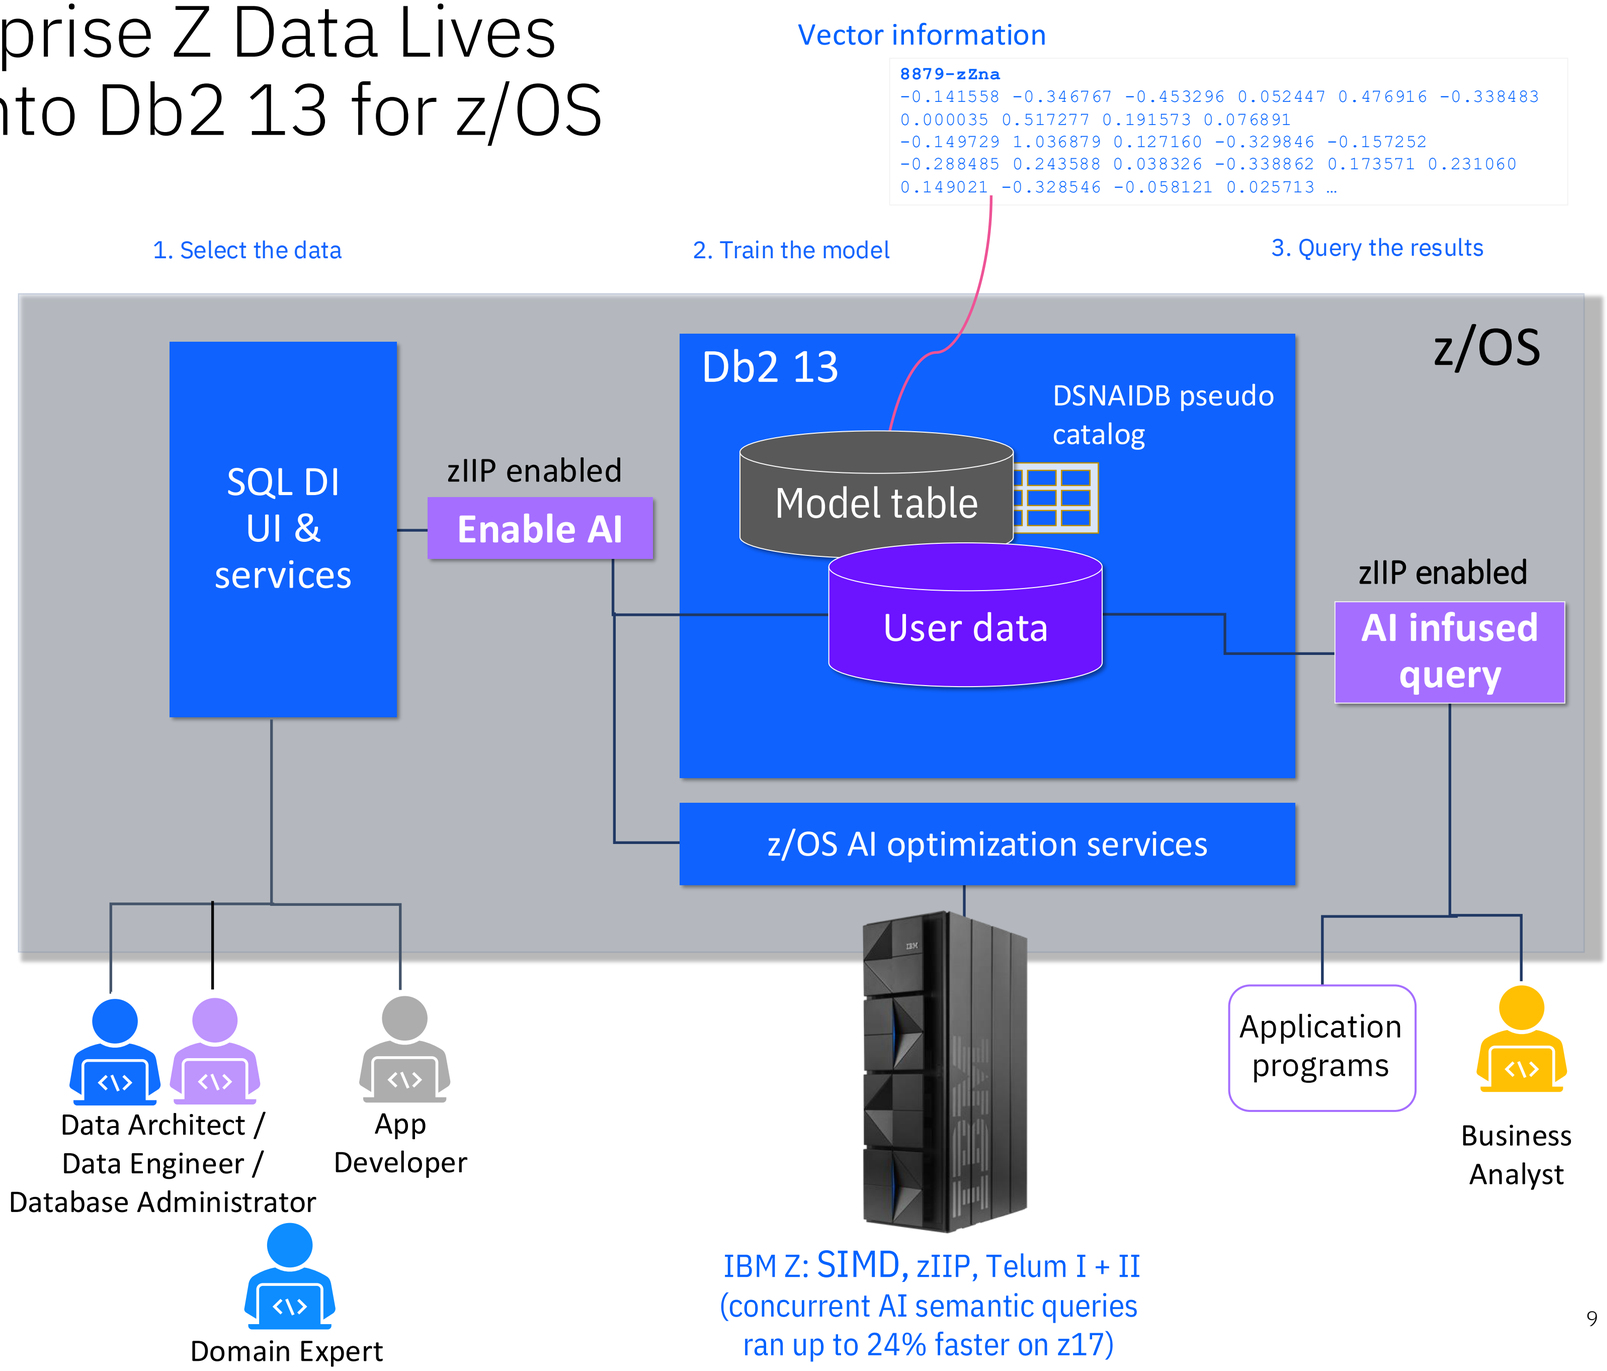

Bringing AI to Where Enterprise Z Data Lives

SQL Data Insights Pro

The Problem

Data gravity and architectural friction

The majority of critical transactional data already resides on Db2 for z/OS. Exporting large volumes of data to external AI stacks introduces latency and cost, also raises concerns around data privacy, regulatory compliance, and governance control.

Shortage of deep AI expertise

Building AI pipelines typically requires specialized skills in data science, model engineering, and infrastructure - capabilities that are scarce and highly competitive to acquire.

The Solution

A **built-in unified embedding model for both structured and unstructured data** enables organizations to search and analyze data by meaning rather than exact matches. AI capabilities are **accessible directly through SQL**, making it easy to integrate into existing workflows. The solution also supports **incremental** training and optimizes model training on IBM Z and LinuxOne, allowing insights to continuously improve as new data arrive.

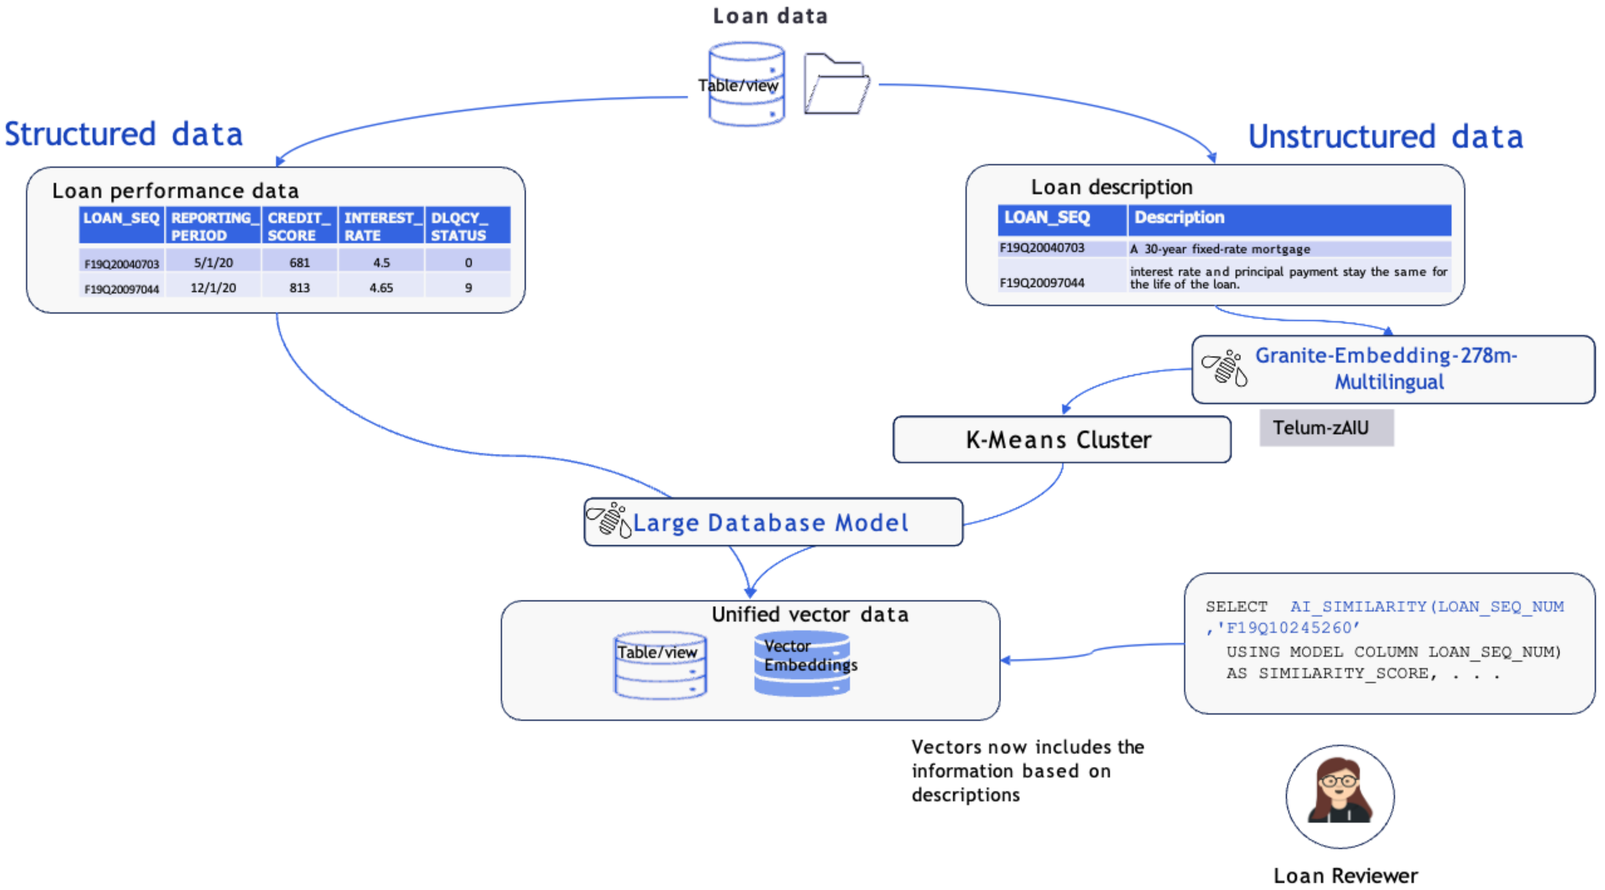

Bringing AI to Where Enterprise Z Data Lives
SQL Data Insights Pro

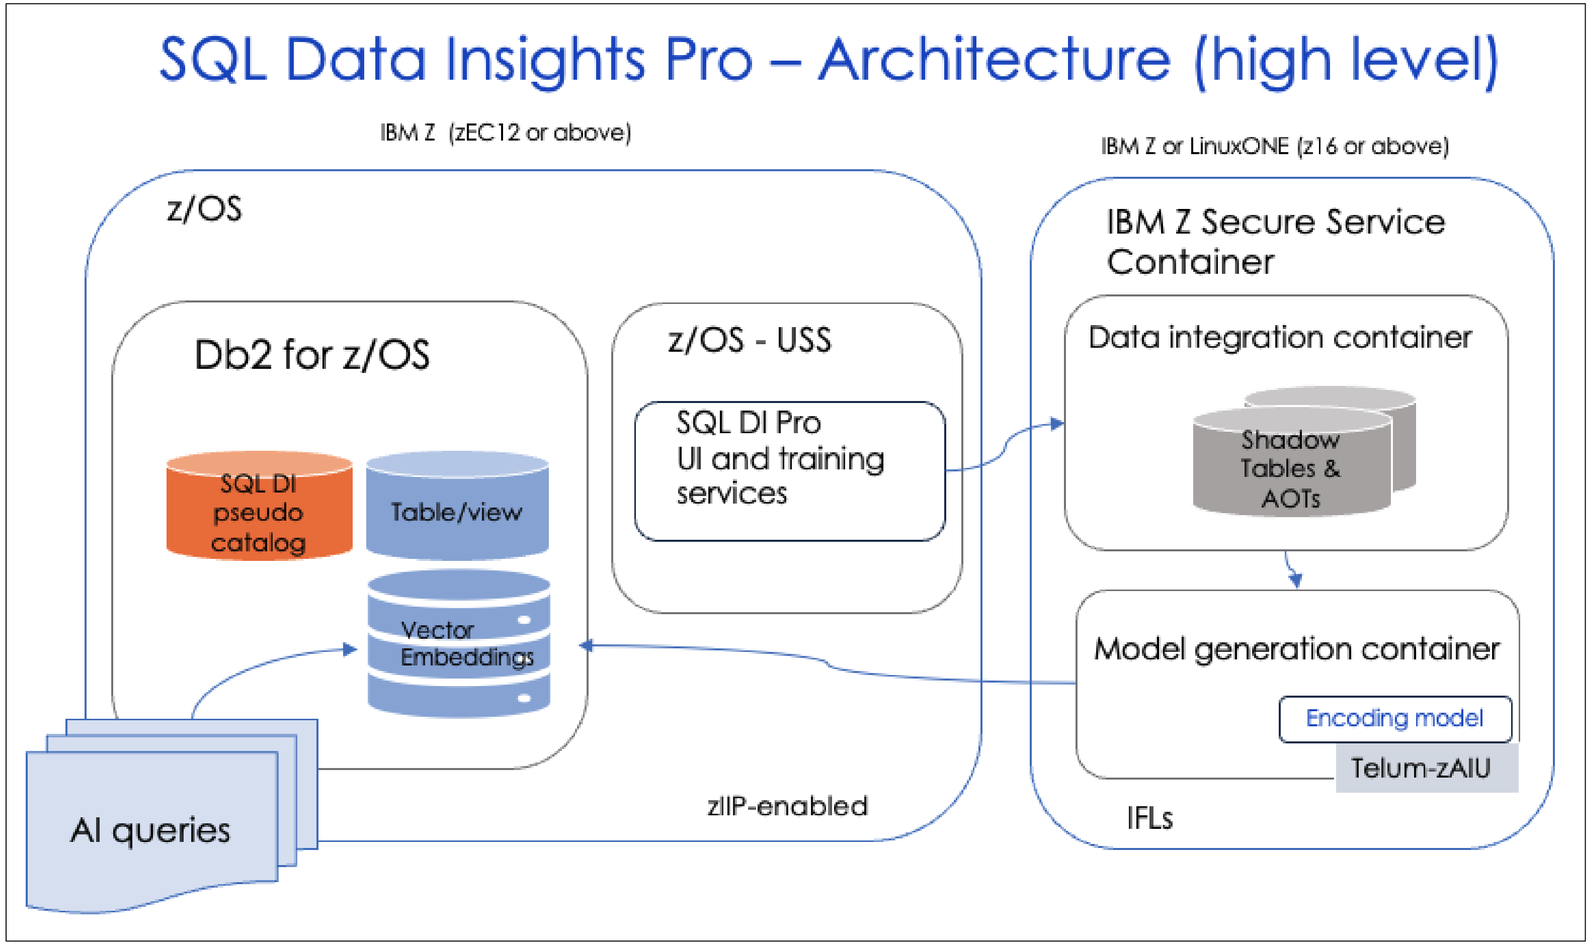

Roadmap

Initial Spyre exploitation

- Accelerate vector embedding (June 2026)
- Natural language to SQL and result summarization
  - AI query/SQL to expand usage beyond SQL users
  - explain why records match or relate - tied back to real business keys and attributes

Feature expansion

- AI query acceleration
- AI Prediction
  - Support unseen data coming from transaction
- AI text query
  - Ask business questions using more natural, sentence-level expressions over text data
- Collaboration with **CAS** to support Hybrid Unstructured files resides both IBM Z and outside

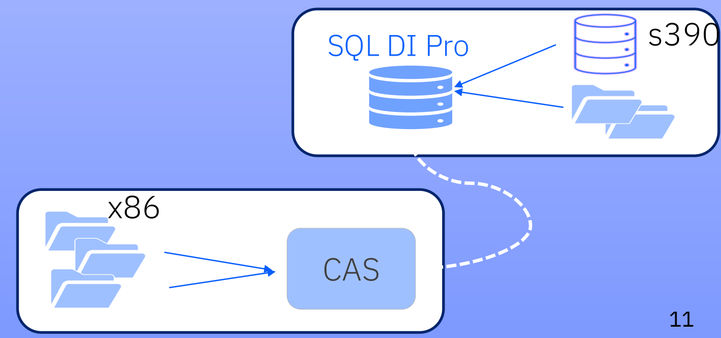

Details
https://www.ibm.com/docs/en/sql-data-insights-pro/1.1.0?topic=overview-sql-di-pro

Don’t move your data … Virtualize it!

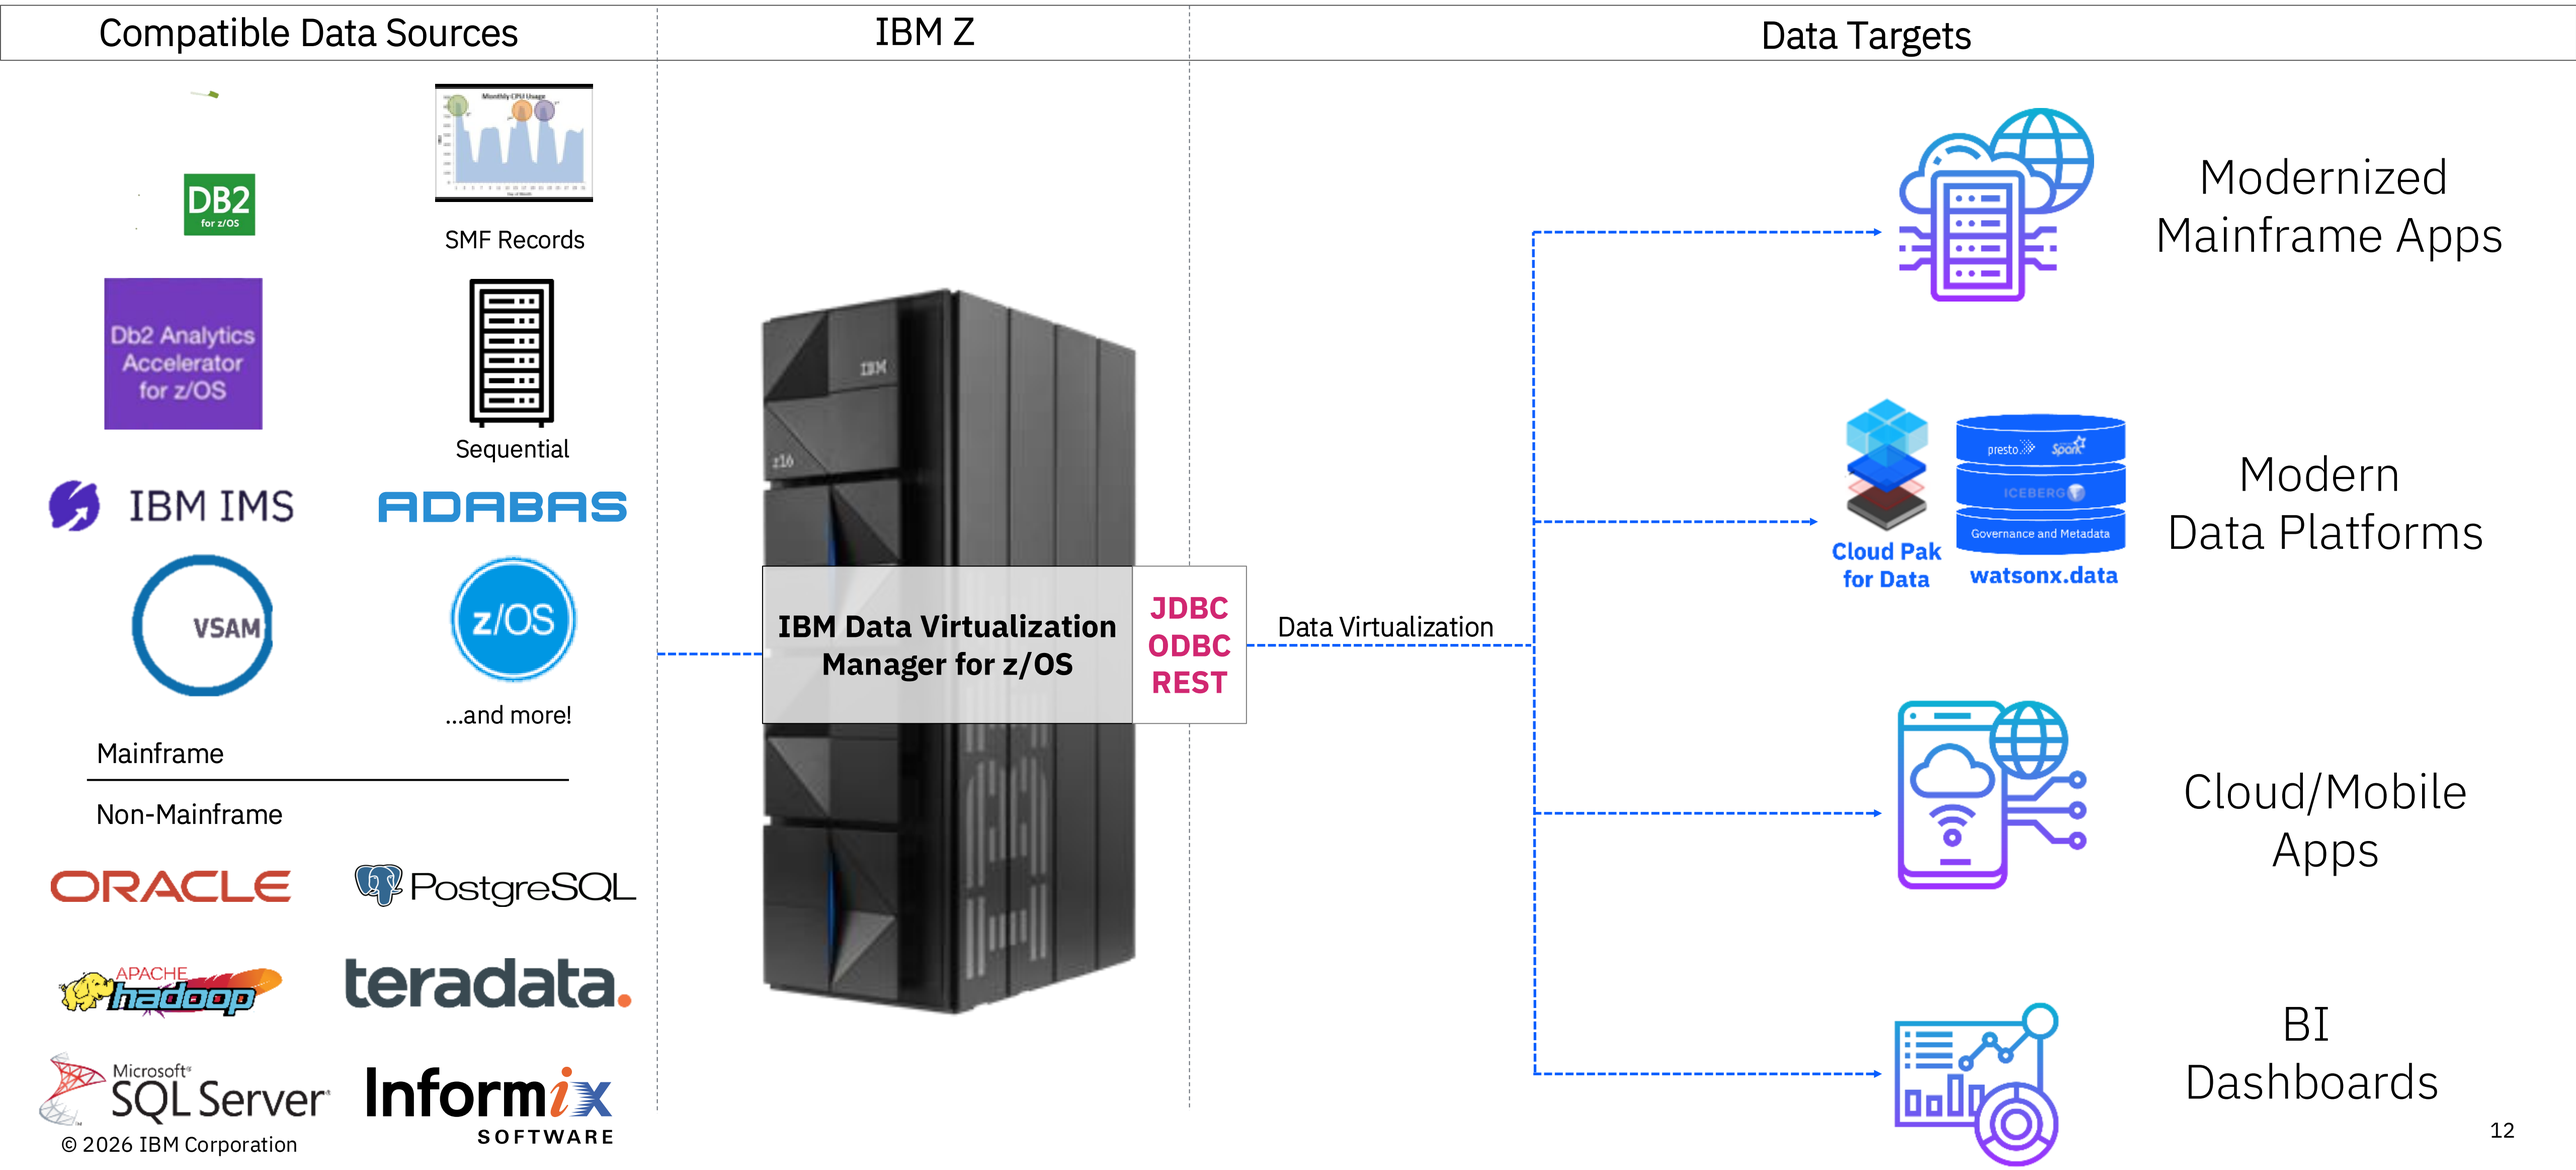

Data Gate: if you need to move Data

- Direct access and virtualization
- Data synchronization / copy

DVM : Data Virtualization Manager for z/OS

The secret sauce: Integrated Synchronization

Simple. Fully integrated into Db2 for z/OS.
Nothing to install on the mainframe.
Easier to maintain.

Isolated. No impact to transaction workloads on IBM Z.

Ultra-efficient. Factors of improved overhead, throughput and latency compared to any other technology.

- Uses ½ the z/OS CPU of traditional replication
- The remaining CPU is +96% zIIP eligible

IBM Z

DWH, ODS, Lakehouse, Downstream Systems

Db2z and/or DVM JDBC

Db2z / IMS / VSAM

Db2 and/or DVM connector

Data Gate

Db2z data *only*

Db2 Warehouse

Db2 connector

Data Gate for watsonx

Db2 OLTP

Db2 connector

Db2z / IMS / VSAM

ICEBERG

Open Data Format on Object Store

Data Gate for Confluent

CONFLUENT

Db2z data *only*

Kafka Connect

Kafka Target

StreamSets

Kafka Consumers

Db2 Data Gate Classic use cases

Transactional caching, Read-only transactions in the Cloud

New digital applications are driving exponential growth in mainframe resource utilization. Provide transactionally consistent systems of record data to cloud-based applications efficiently and securely.

Data Fabric

Make data available and synchronized for ready access within a data fabric, including one-click integration with Watson Knowledge Catalog. Simplify the integration of data sources within your data fabric.

Db2 Efficiency and Archiving

Redirect Db2 for z/OS analytical queries to Db2 Data Gate without consuming mainframe resources and without impacting operational SLAs.
Archives Db2 data to a Db2 Warehouse database on Cloud Pak for Data.

Z Lakehouse - Data Gate for watsonx use cases

Efficiently integrate mainframe data in a data lakehouse for Analytics & AI

Unleash the power of IBM Z data for analytics and AI in a lakehouse with efficient integration to via IBM Data Gate for watsonx.

Integrate your critical IBM Z data with Salesforce

Enables you to activate AI agents within Salesforce, offering improved, personalized, customer interactions with easy integration of IBM Z data into the Salesforce Data Cloud.

Augment your existing cloud warehouse/lake

Reduce proprietary data ingestion costs by utilizing the Iceberg lakehouse architecture. Replace expensive data pipelines with a more efficient strategy.

Future proof your data

Replicate once – use many… Utilize an open lakehouse architecture to minimize the number of copies of enterprise mainframe data sent to the cloud. One (efficient) copy can be used for multiple products and use cases.

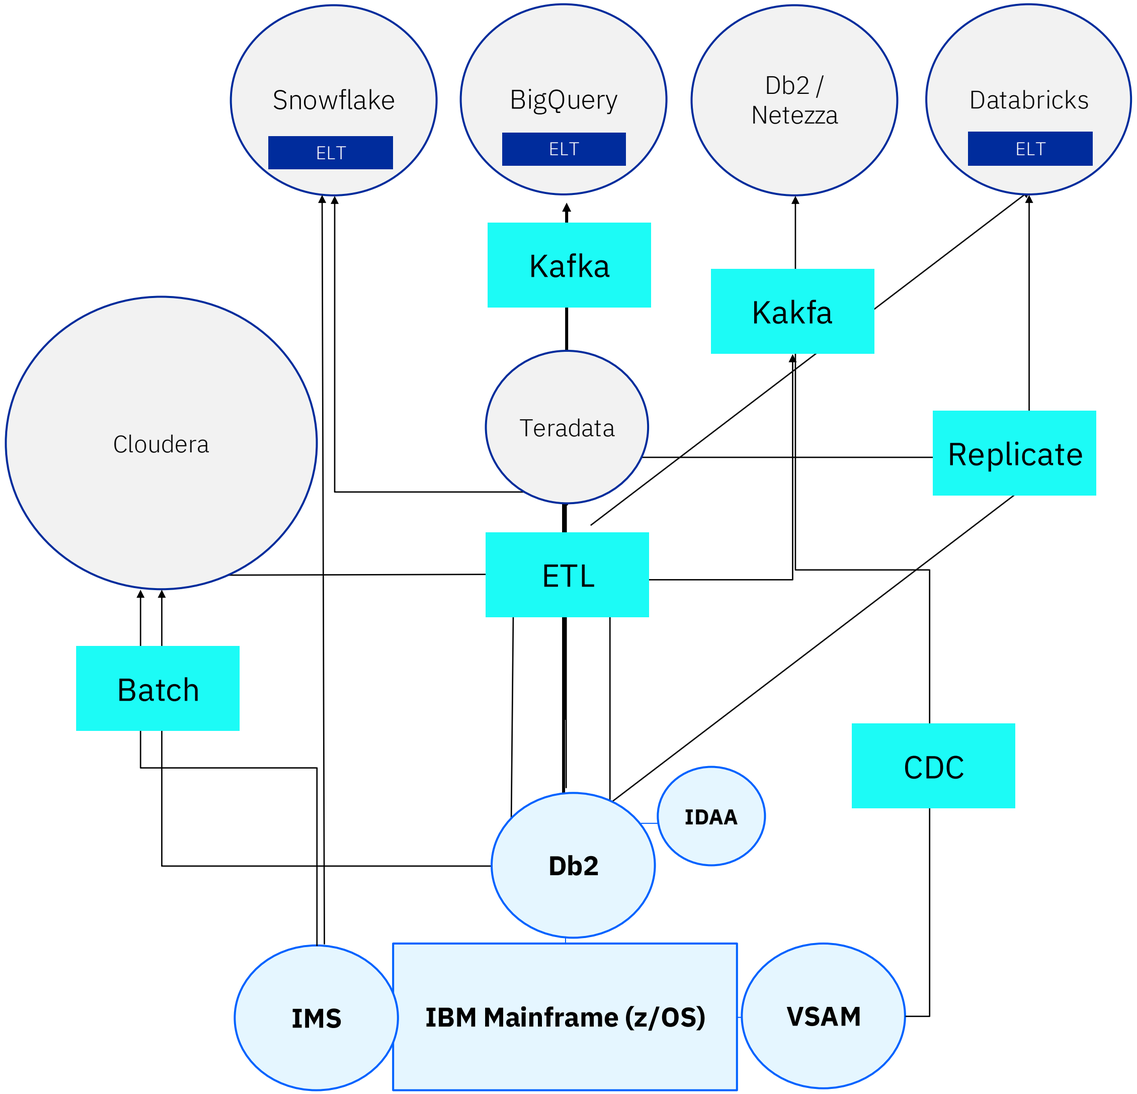

From this...

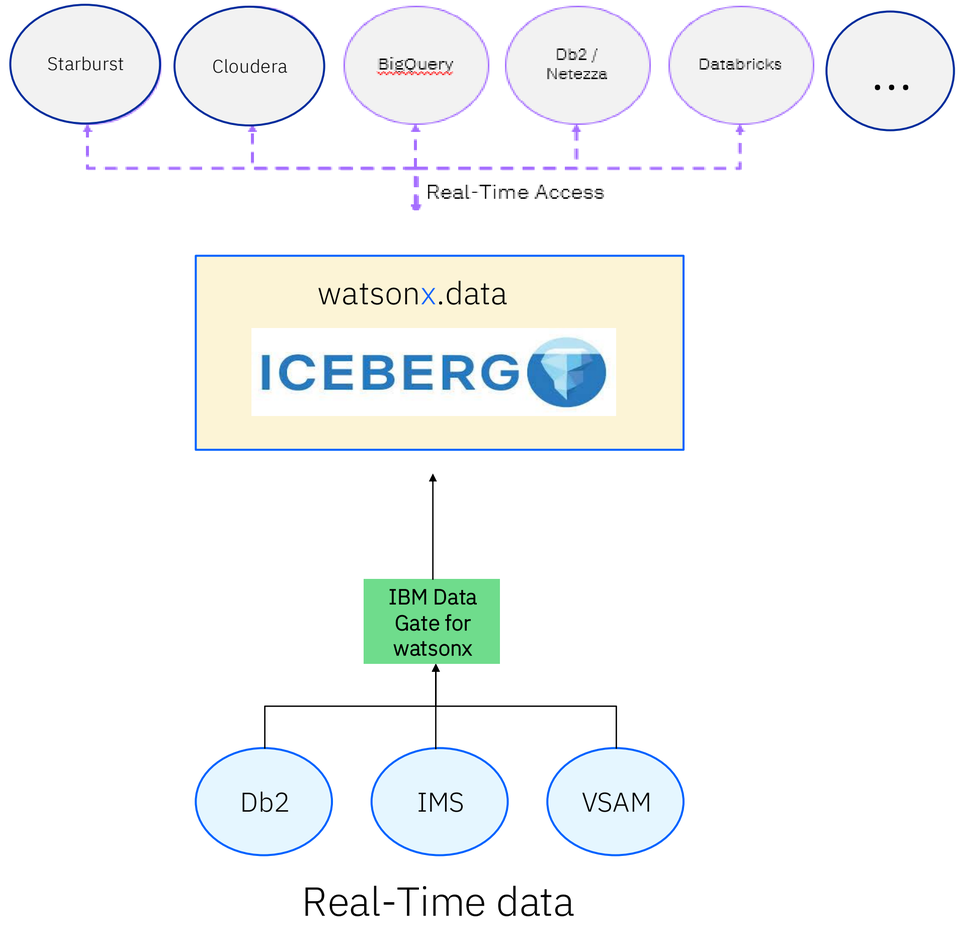

To this!

BANCO DO BRASIL

Banco do Brasil turns data into powerful insights

“The bank’s branches take advantage of over 70,000 customized, user-built, reporting dashboards that depend on fast access to the most current mainframe data provided by the Db2 Analytic Accelerator.”

Eduardo Pingarilho Mendizabal
Mainframe and Database Manager

The Challenge

Banco do Brasil needed faster, lower-cost access to IBM Z data to support over 70,000+ dashboards. Slow, expensive reporting (24+ hours) delayed managerial insights and productivity tracking.

Read the full story

The Solution

IBM Db2 Analytics Accelerator

The Impact

- Fully integrated with Db2 for z/OS to offload analytics, reduce CPU use, and deliver real-time mainframe data for faster reporting and decisions.
- 100,000s of employees gain real-time insights from complex queries on live mainframe data.
- Loan-processing data time cut from 24 hours to minutes.

Garanti BBVA International

Accelerate analytics deliver banking insights

“As a major financial company, we have a team monitoring performance of the database and application environment. With the IBM Db2 Analytics Accelerator, we can identify most problems before they impact the business.”

Semra Arslan, Manager
Z Cloud Data Technologies, Garanti BBVA

The Challenge

Garanti BBVA faced billions of daily transactions, heavy batch workloads consuming CPU and risking core banking performance, and slow analytics on static batch data—driving the need for secure, cost-efficient, real-time insights without overloading IBM Z or disrupting core banking.

The Solution

IBM Db2 Analytics Accelerator

Read the full story

The Impact

- Enabled to run analytics and compliance reporting 3000X faster.
- Saved 45 mainframe CPU hours per day.
- Reduced key regulatory report time from 2 days to 1 minute.
- Accelerated 300+ nightly batch jobs and analytics across 33 TB + 44 TB of data respectively.

وزارة الداخليـــــــة
Ministry of Interior
State of Qatar • دولـــة قطـــر

Lightning-fast query processing and lower cost

“ Our Db2 for z/OS database is capable of doing everything. ”

Azeem Mohammed
Database Administrator for IBM DB2 for z/OS
Ministry of InteriorState of Qatar

The Challenge

Rising Db2 for z/OS workloads drove CPU and memory spikes, slowing dashboards and reports and increasing costs. Long batch runs and data movement workarounds blocked timely analytics. The Ministry needed workload offload and fast, reliable access to current transactional data.

The Solution

IBM Db2 Analytics Accelerator
IBM Db2 for z/OS

Read the full story

The Impact

— Offloaded heavy Db2 analytics, reducing latency, lowering CPU cost and enabling near–real-time reporting.
— 98% reduction in latency — insights delivered in seconds from minutes.
— Reports 30X faster (20 minutes → 40 seconds).
— Data sync improved from 3 minutes to ~3 seconds, reducing CPU consumption.

die Mobiliar

“ We worked together with IBM to process insurance offer recommendations faster and more accurately. By unlocking hidden data patterns with NLP-based AI functions in near-real time while ensuring privacy and security, we saw an 94% accuracy in prediction results. These very promising results have motivated us to integrate this technology into our underwriting processes in the near future.”

Thomas Baumann
Data Evangelist/ IT Architect
Swiss Mobiliar

The Challenge

The customer wanted to predict the probability that an auto insurance offer would be converted into a contract immediately after the offer was created.

The Solution

IBM Db2 for z/OS with IBM SQL Data Insights

Read the full story

The Impact

— Client achieved near-real-time AI inferencing with 94% prediction accuracy.
— Ideation to production time reduced to 4 hours.

Questions ?

Bring AI Where It Matters

AI and analytics execute where they deliver the most value — in place, on platform, or in the Hybrid cloud. The right execution location depends on the workload, latency requirements, security requirements, data size, governance posture, and scale of the use case.

Embedded AI Insight

Business Scenarios

- Investigating suspicious transactions using semantic similarity
- Exploring related claims, accounts, or policies
- Pattern discovery and decision support on live operational data

Technology Fit

SQL Data Insights & Pro

- Embeds AI directly inside Db2
- Generates and queries vector embeddings in place
- Preserves governance and security boundaries
- Eliminates replication and ETL

When to Choose This

- Insight must operate on fresh system-of-record data
- Data cannot move due to risk or compliance
- The objective is decision support, not large-scale model training
- Analysts or operational workflows need AI-enhanced exploration
- Latency and governance matter more than elastic compute

This is AI embedded where the data lives.

IBM Z Core

Accelerated Analytics on Platform

Business Scenarios

- 360° customer view
- Regulatory and compliance reporting
- Loan portfolio and risk aggregation
- High-volume analytic dashboards
- Cross-domain reporting across Db2 subsystems

Technology Fit

Db2 Analytics Accelerator (IDAA)

- Delivers high-speed analytic query acceleration
- Isolates analytics from OLTP workloads
- Eliminates ETL/ELT complexity
- Keeps data on Z under Db2 governance
- Enables near real-time HTAP

When to Choose This

- Workloads require heavy aggregation or complex joins
- Analytics must not impact transactional performance
- Near real-time is sufficient (not inline workflow AI)
- Data should remain on Z
- Enterprise reporting and BI acceleration are priorities

This expands intelligence on platform without moving data outward.

Linux environment on IBM Z (IFLs)

IBM Z with Virtualized Access

Business Scenarios

- Exposing Z data through REST APIs, JDBC, ODBC
- Supporting hybrid cloud applications
- Enabling data mesh participation
- Integrating Z and non-Z data sources
- Modernizing legacy applications

Technology Fit

Data Virtualization Manager (DVM)

Provides real-time federated access

- Avoids full data replication
- Enables modern SQL and API access
- Reduces architectural complexity
- Supports hybrid app development

When to Choose This

- Applications need current data,
- The goal is modernization and direct access
- Analytics is secondary to data availability but can be combined with IDAA if analytics is required.
- Hybrid architectures require integration
- Governance must remain centralized

This enables modernization without unnecessary data movement.

Z Virtualized
IBM Z or Linux env on IBM Z

Cloud & Hybrid Environments

Business Scenarios

- Large-scale AI model training
- Enterprise data science initiatives
- Integration with Snowflake, Databricks, Salesforce
- Cross-domain AI workloads
- Lakehouse-scale analytics

Technology Fit

Data Gate (Classic & for watsonx.data)

This technology:

- Synchronizes selected Z data to open Iceberg tables
- Offloads workload from the core system
- Enables elastic compute environments
- Maintains governed and consistent replication
- Supports enterprise-scale AI platforms

When to Choose This

- AI spans multiple enterprise domains
- Elastic cloud compute is required
- Cross-platform data fusion is needed
- Latency is not transactional
- Replication is acceptable and governed

This enables enterprise AI beyond the core platform.

Hybrid / Cloud

System of Record

On-Platform Expansion

Elastic Enterprise AI

Thank you