zData Integration Strategy

M. Cüneyt Göksu
Executive IT Specialist
IBM Germany
Mehmet.Goksu@ibm.com

11 – 13 May 2026

IBM

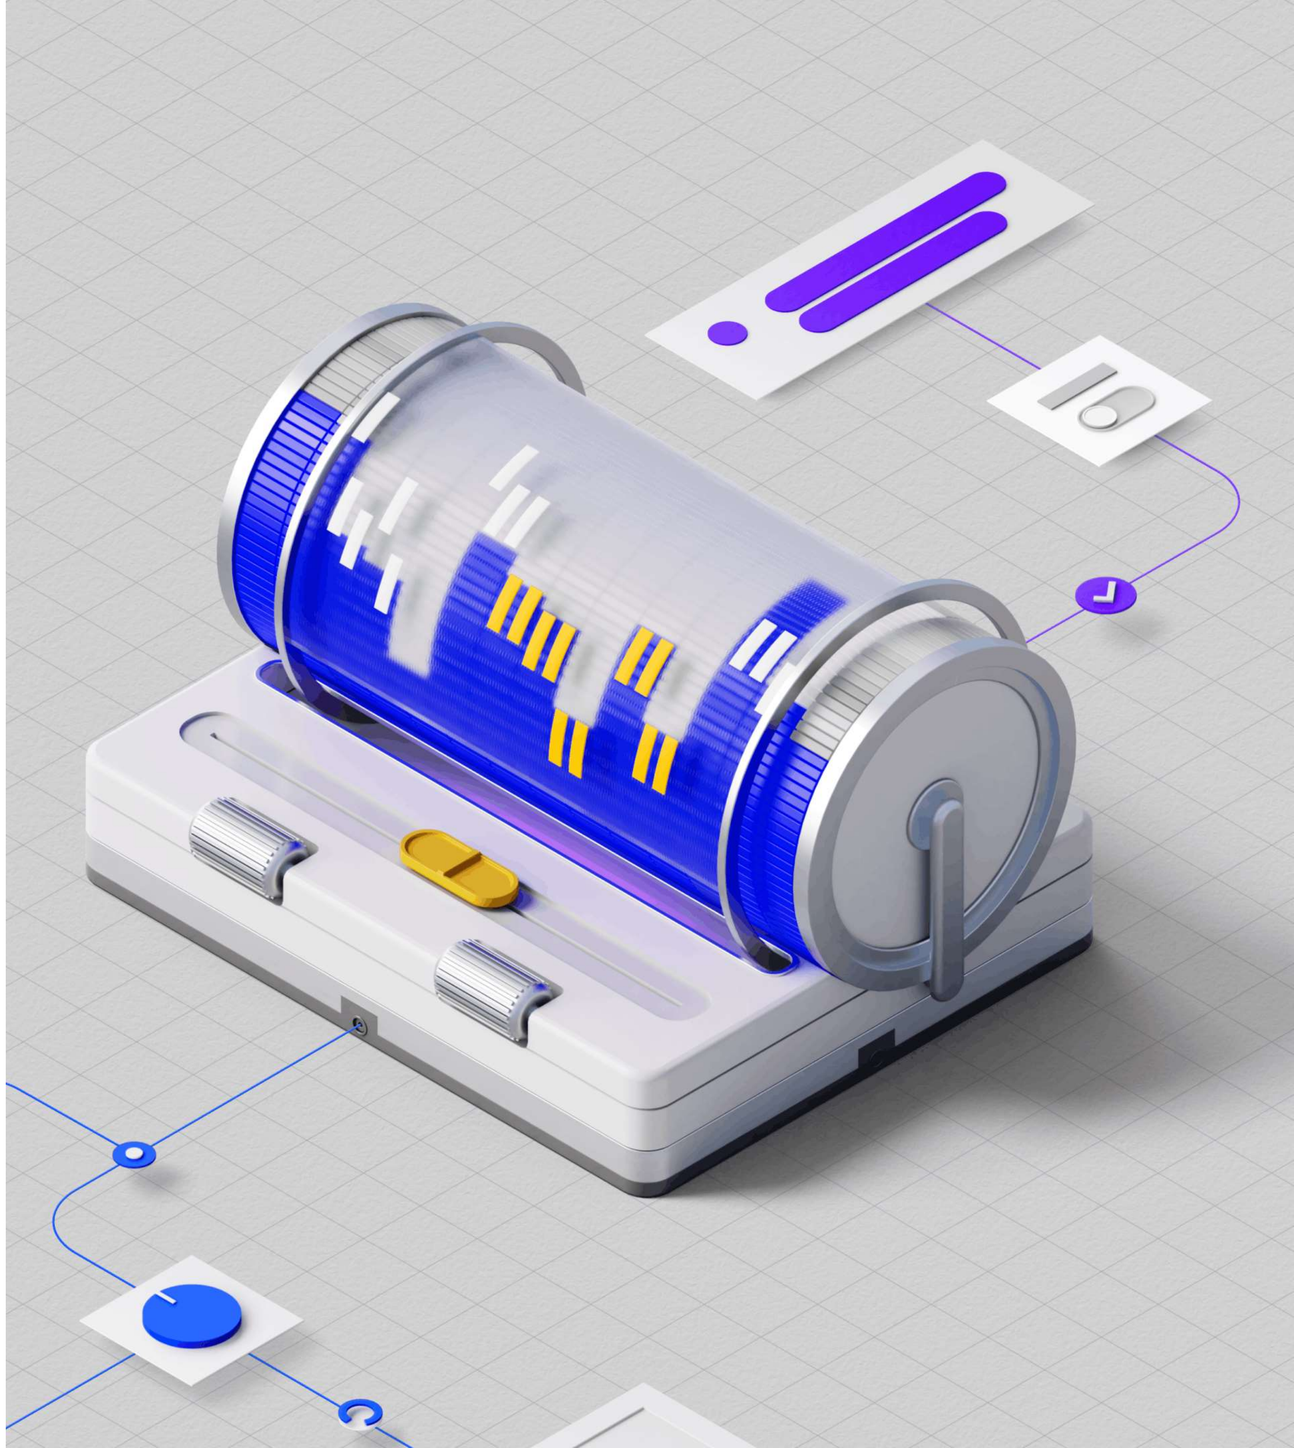

Agenda

Status

The IBM Z Data Value Pyramid

Architecture

Unlock real-time & secure insights from the high value transactional Z data

70%

of all global transactions by value run on a mainframe[1]

79%

of executives agree that mainframes are essential for enabling AI-driven innovation[2]

91%

of top-performing companies invest in AI and data[3]

1.Source : Mitigating Fraud in The AI Age: Supporting Transaction Fraud Detection at Scale on IBM z17, Celent, April 2025
2.Source: Mainframes as mainstays of digital transformation, IBM Institute for Business Value, October 2024.

Business innovation requires a thoughtful Z data strategy

Security & risks in data movement
Organizations are hesitant to move data from IBM Z due to possible risks, strict security, regulatory, and compliance requirements

High-cost movement of transactional data
Moving data off IBM Z for analytics and AI increases cost, latency, and operational complexity

Insights not based on real-time data
Batch-oriented data pipelines result in analytics and AI insights derived from stale rather than live transactional data

AI & Analytics is only as good as your **data**

Location Matters for AI & Analytics Decisions

AI & Analytics Off-Platform

Off-platform AI creates distance between data and decision, adding:

- X Latency penalty
- X Compliance risk exposure
- X Infrastructure complexity
- X Competitive disadvantage
- X Huge movement costs

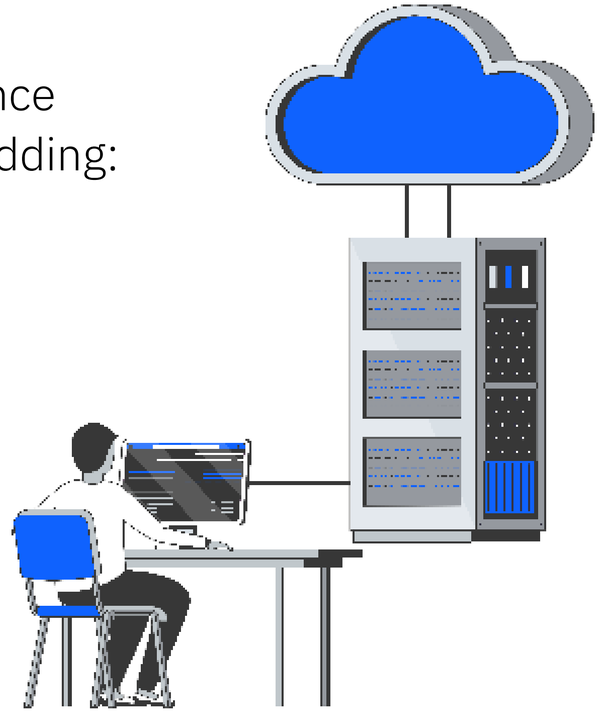

AI and Analytics on IBM Z

Run AI inferencing on the platform where mission-critical data lives and decisions actually take place for co-location benefits:

- ✓ Real-time speed advantage
- ✓ Security advantage, no transfer risk
- ✓ Significantly lowered costs possible

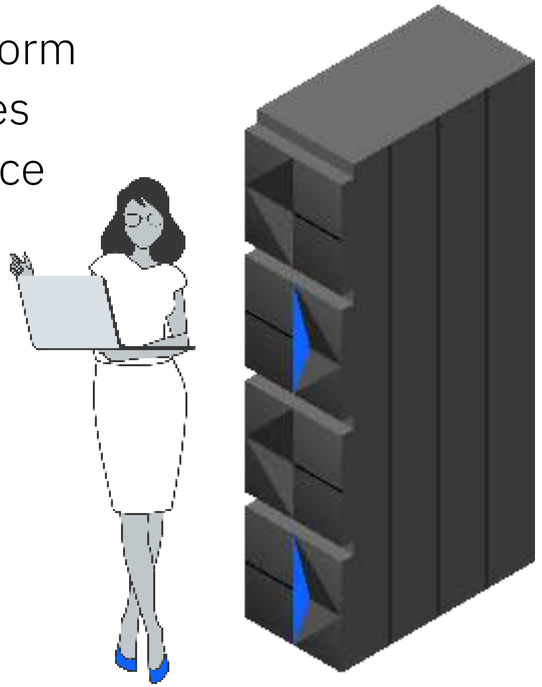

Agenda

Status

The IBM Z Data Value Pyramid

Architecture

Unlocking real-time AI value from applications and the most trusted data

IBM’s approach builds value “**bottom up**” without disturbing what already works

— Leverage AI to **magnify the value of real-time transaction processing** applications - efficiently, securely, and at scale.
— Enable companies to **deploy AI agents, digital workers, and autonomous workflows** with live—not static data.
— Boost productivity and operating efficiency with **AI that delivers actionable insights** and automates data handling at scale.

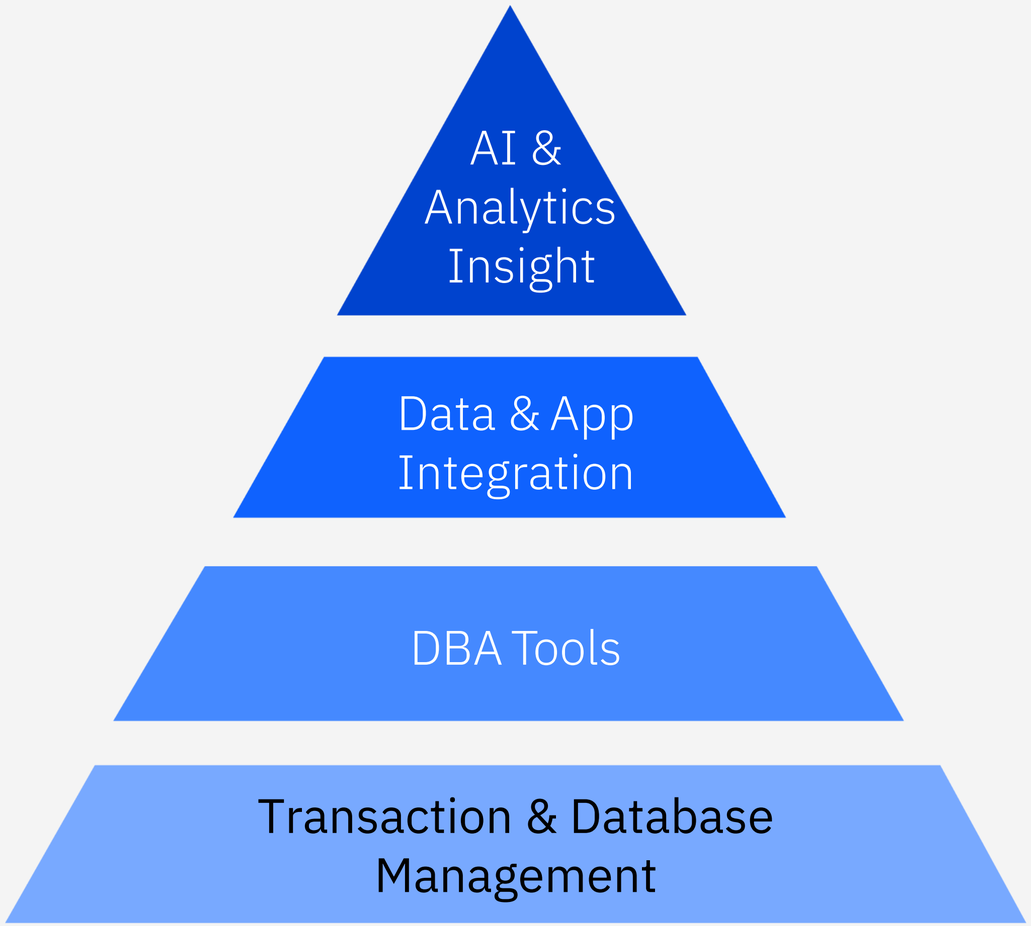

Powerful database foundation to make Z data agentic-ready

- IBM Z is the source of truth for the business
- Preserve the integrity of the system of record while unlocking greater value from it.

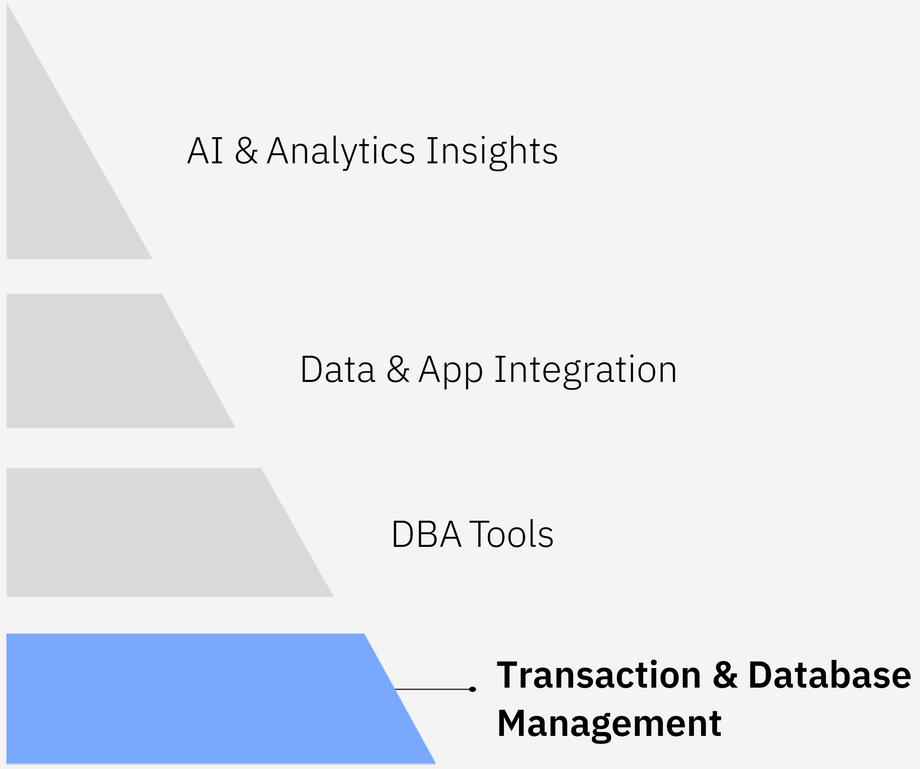

AI-Driven Intelligence
Infuse AI into Z data to automate insights, optimize performance, and enable real-time decisions.

High Performance Access
Enable high-performance access to Z data, allowing agents and apps to interact with transactions and data in real time.

Enterprise-Scale Operations
Enterprise scale design for massive concurrency and high throughput on critical workloads.

Trusted & Governed Interactions
Provides secure access to Db2 data to leverage industry-leading security, governance, and compliance.

IBM’s Db2 and IMS tools portfolios are designed to be force multipliers for the DBA, giving them ***a comprehensive, integrated toolkit that simplifies administration, enhances system health, and reduces the operational burden of managing mission critical databases.***

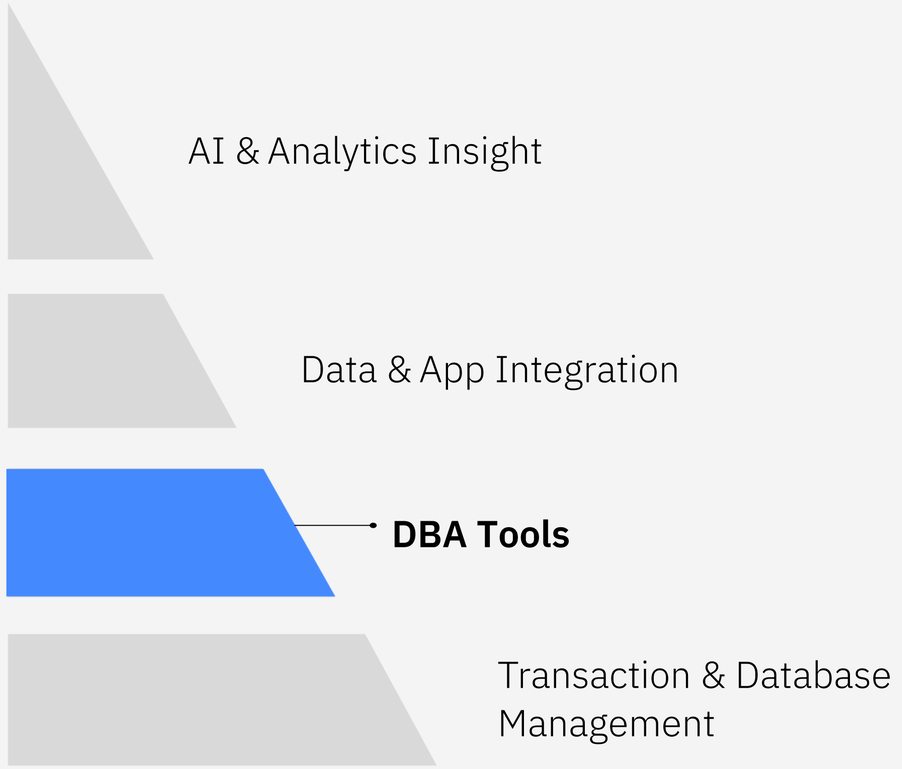

Agent-driven performance optimization
Automatically tunes Db2 and IMS workloads to improve response times and resource efficiency.

Automated DBA operations
Db2 AI for z/OS provides machine learning driven SQL optimization, anomaly detection, system assessment, and distributed connection control, allowing DBAs to offload repetitive tuning tasks

Predictive issue prevention
The Db2 Tools for z/OS suite streamline

- Database Maintenance,
- Performance Tuning,
- Schema Management,
- Utilities Automation,
- DevOps Integration
- Recovery & Restart

Accelerated modernization
zDBA—a DBA persona centric, agentic experience for Db2 and IMS tools built on the watsonx Assistant for Z (WXAZ) platform. .

The Challenge of:

DBA Operational Overload

Reactive Firefighting

Manual, immediate responses to critical alerts, performance degradation, and outages.

DevOps

Review database-related change requests or schema changes submitted by development teams for eventual deployment.

Performance and Health Monitoring

Maintaining optimal performance and spotting issues before they impact users.

Routine Maintenance Chores

Scheduled, repetitive tasks like running index rebuilds, checking fragmentation, log backups, provisioning, de-provisioning, and manually troubleshooting connectivity/permissions.

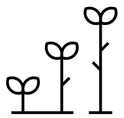

Proactive Skill Evolution

Staying current with skills while coaching, training, and mentoring junior colleagues.

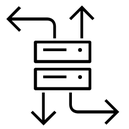

Industry Trend

Multi-Database Dexterity

Skill in managing different databases.

Introducing the solution…

IBM Z Database Assistant
the DBA Experience Re-Imagined

…a unified, intelligent workspace creating a force multiplier for every DBA.

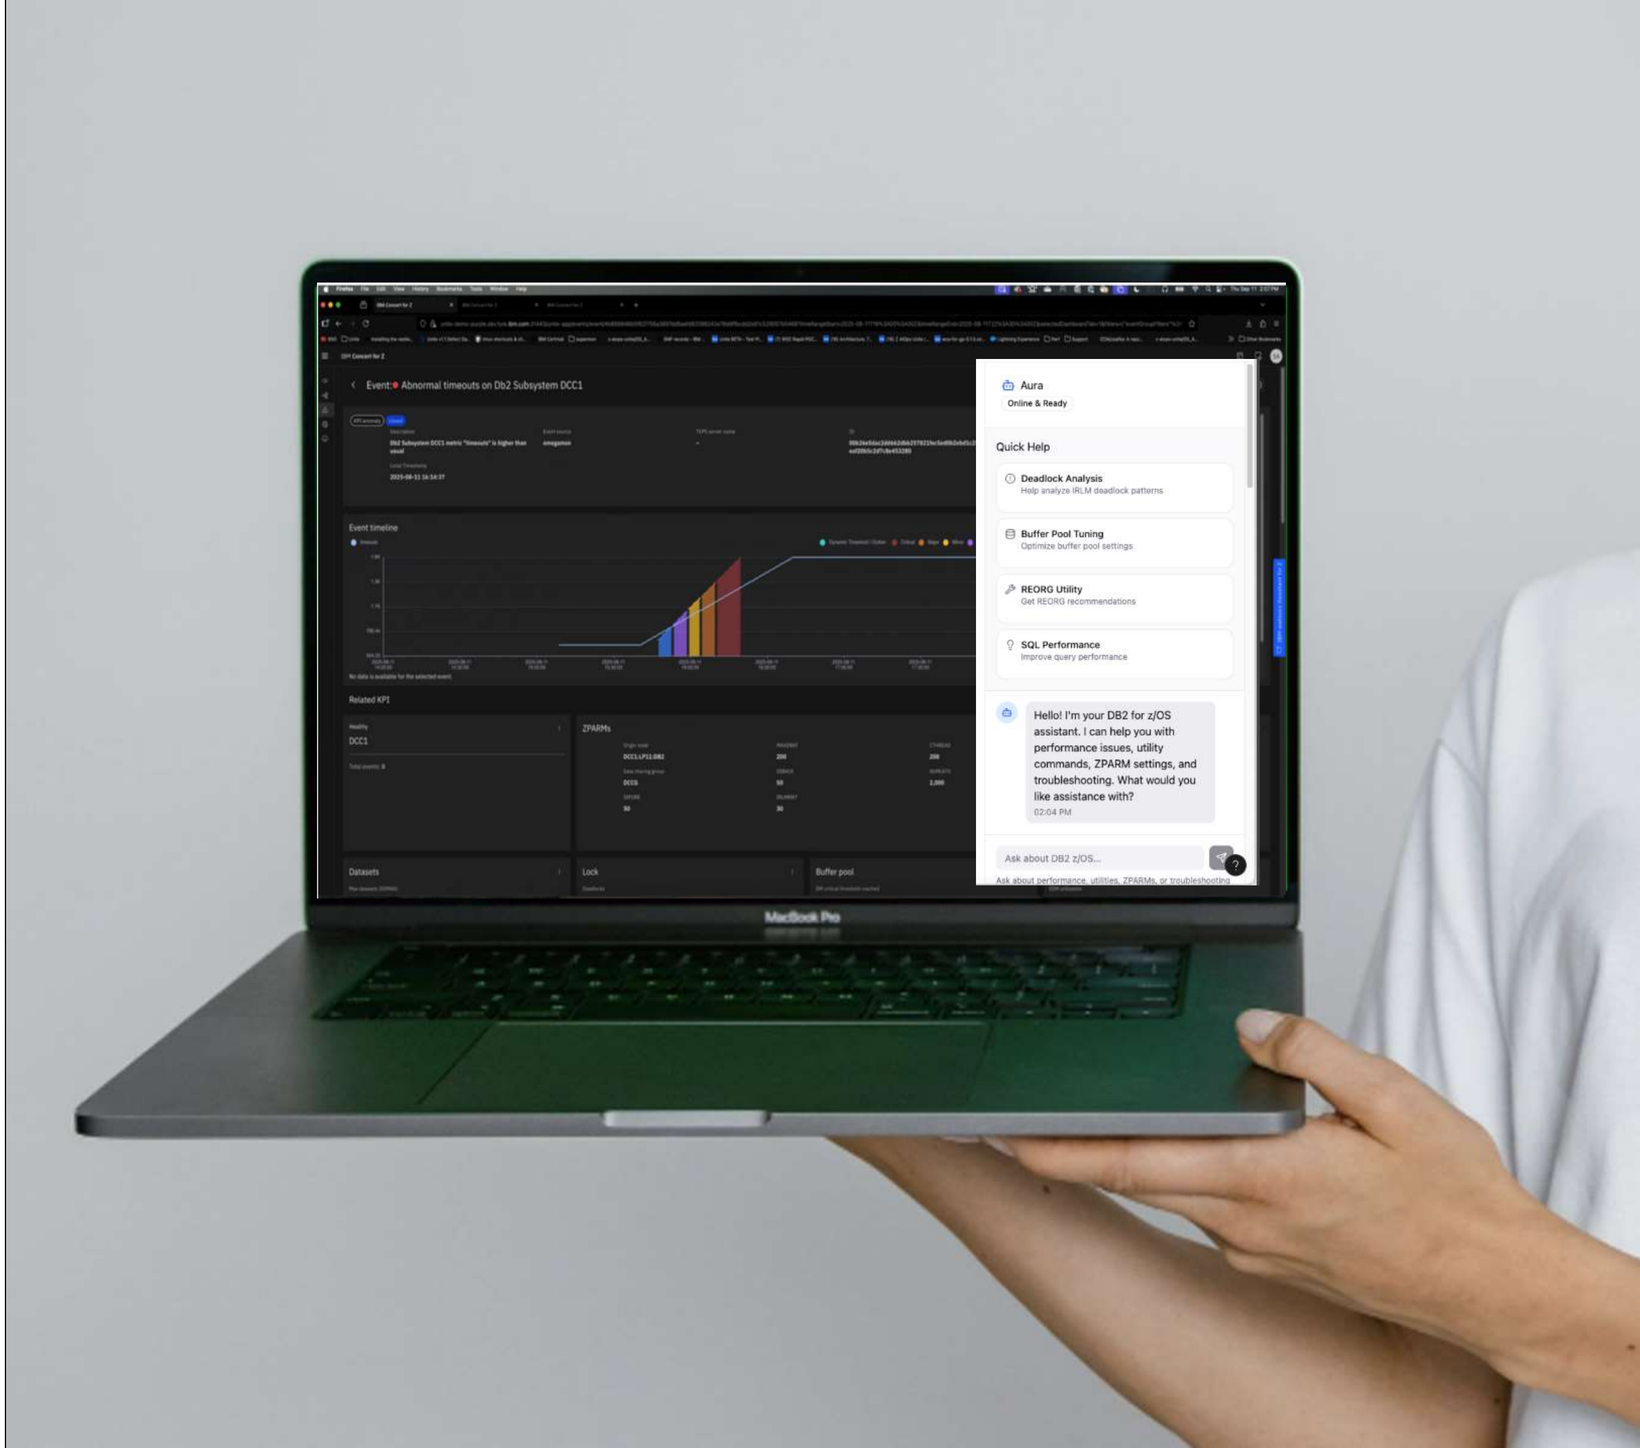

Value Propositions

Continuous Observation

Monitors Db2 and IMS environments 24/7.

Plain Language Explanation

Translates complex issues into simple, understandable terms.

Actionable Recommendations

Provides clear, effective steps to address issues.

Safe Execution

Orchestrates and executes actions using existing Db2 and IMS core and tools features.

Efficient data and application integration

This layer is the bridge between **systems of record** and **systems of insight**. It enables real-time access to trusted Z data across applications, analytics platforms, and AI workflows — **without impacting core systems**

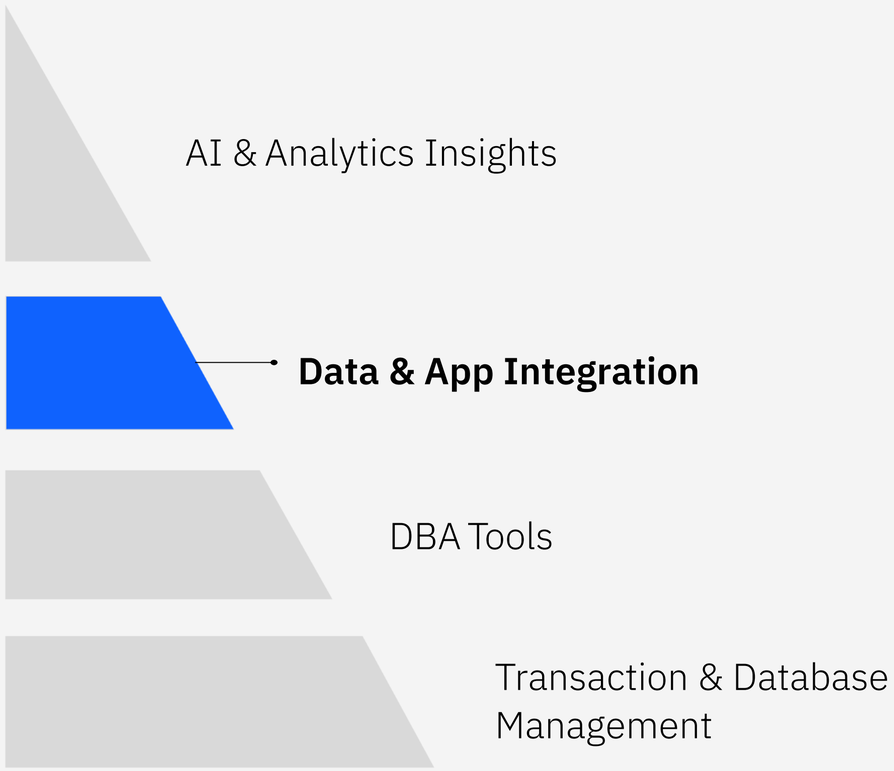

Unified Data Foundation
Securely connect Z data to lakehouse for faster AI training and real-time insights.

IBM Z Digital Integration Hub (zDIH)

Data Virtualization (DVM)

Accelerated Innovation & Time-to-Value
Embed AI into transactional workflows, uncover hidden patterns, and automate decisions through secure API integration.

SQL Data Insights (SDI)

Simplified Hybrid Cloud Integration
Seamlessly connect IBM Z with cloud-native apps to cut complexity, cost, and latency—while ensuring governance and compliance.

IBM Data Gate

Enhanced Cost Efficiency
Reduce data movement and infrastructure costs while keeping Z data secure—enabling AI and analytics without compliance risk.

The Db2 Analytics Accelerator (IDAA)

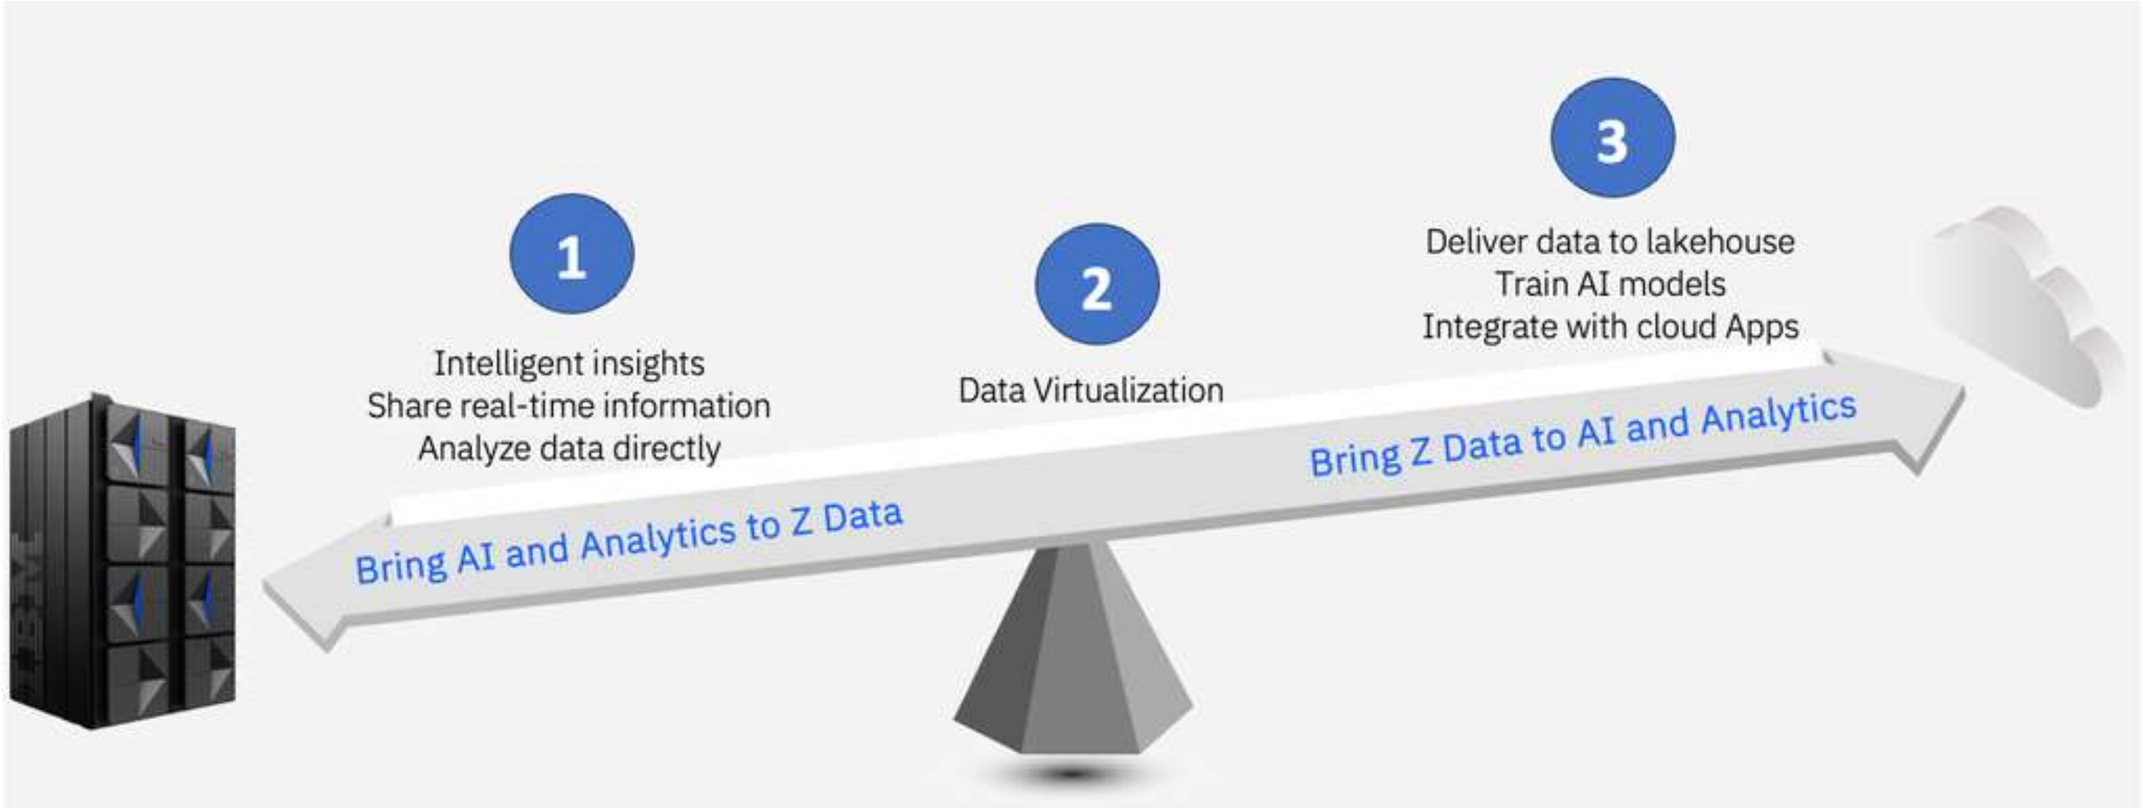

1. Bring Analytics & AI to the Data on the mainframe
   When freshness, security, and low latency matter, analytics run directly on Z.
2. Keep Data on the mainframe and Expose It Virtually
   Virtualization delivers current data to modern applications without replicationg the system of record.
3. Selectively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

Agenda

Status

The IBM Z Data Value Pyramid

Architecture

IBM Z Data Architecture

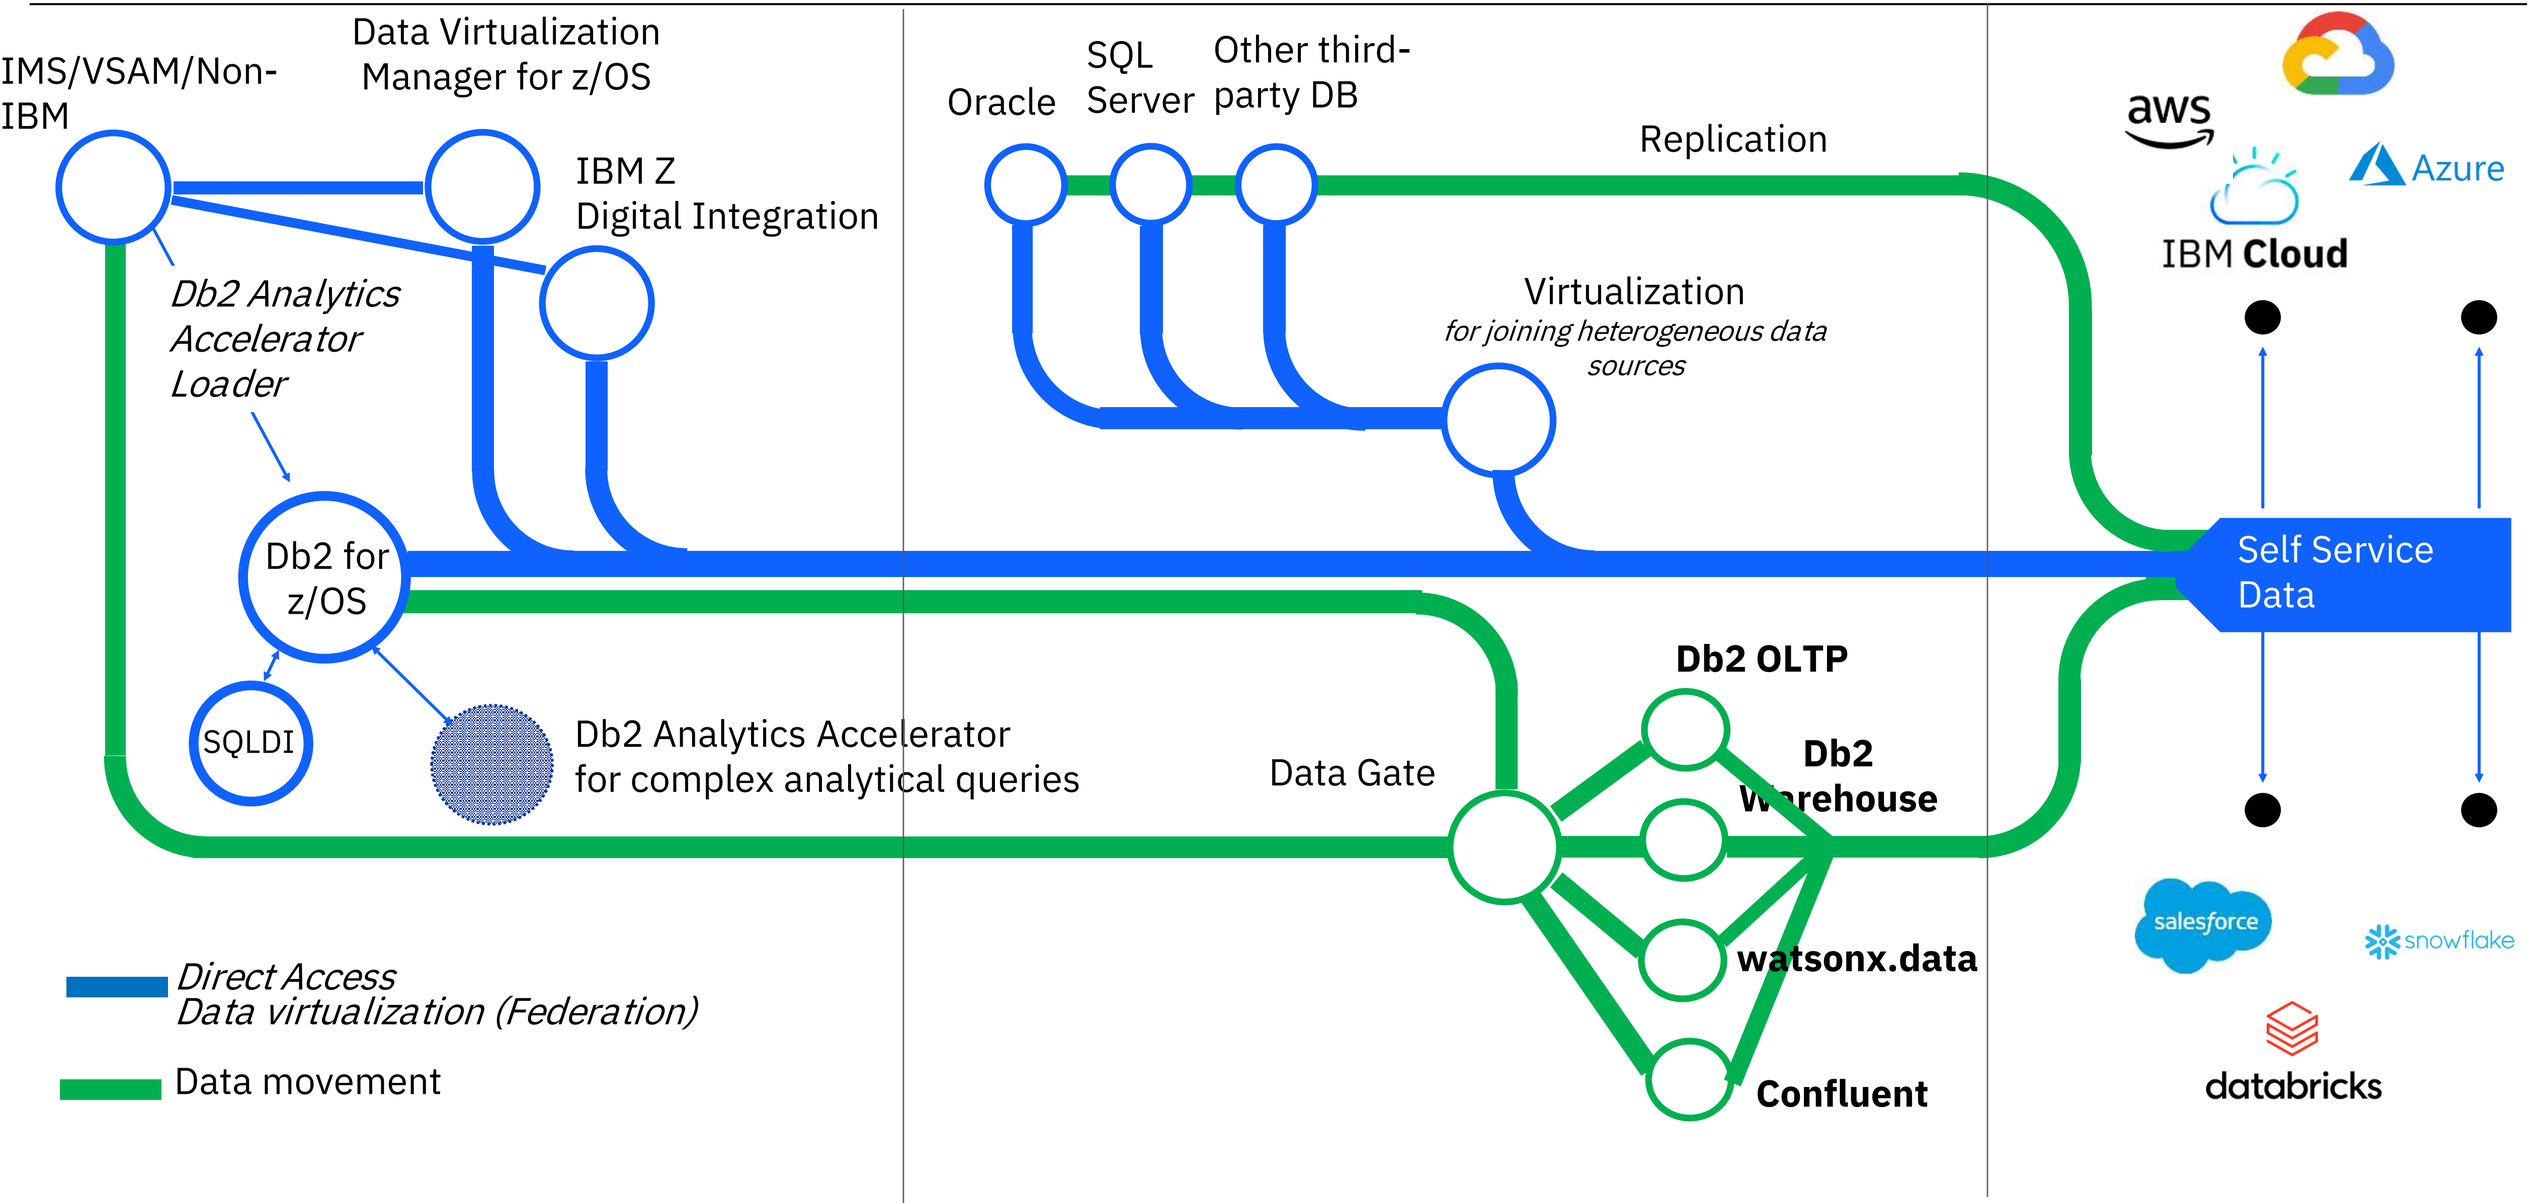

Hybrid Analytical and Transactional Processing - Db2 for z/OS and IDAA

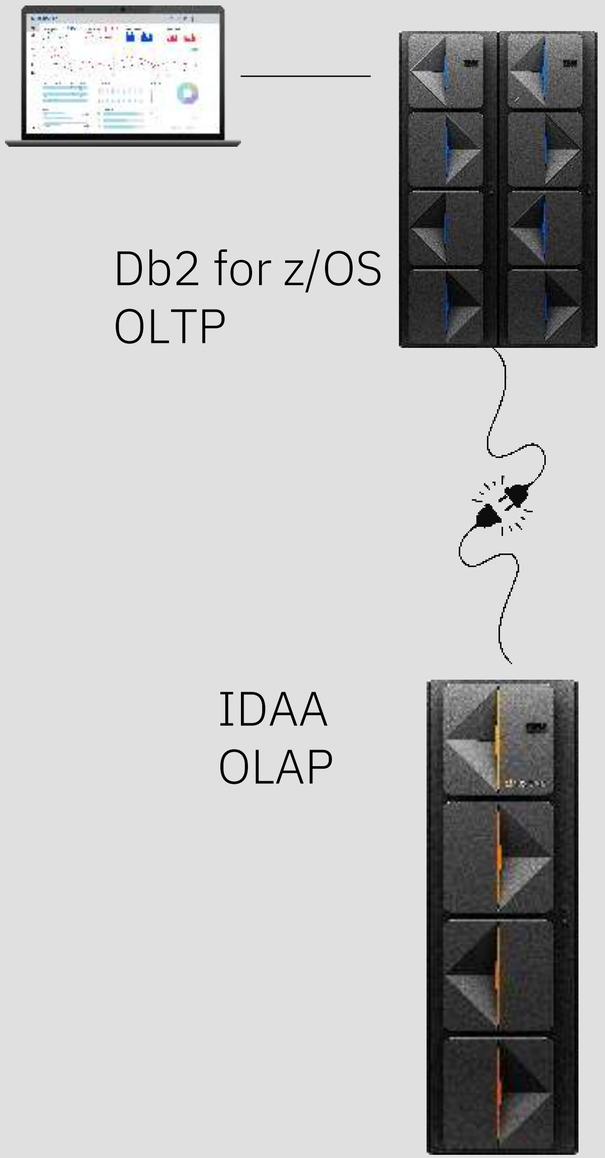

Ultimate performance

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

Performance Warehouse

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

Db2 13 with SQL Data Insights included

Uncovers and monetize hidden insights in your most important data

- Db2 builds an unsupervised, Large Database Model against selected Db2 data, to answer business questions
- Provides the ability to discover patterns in their rich operational data and easily answer business questions
- Opens AI to the business through familiar SQL, without requiring AI expertise or a Data Scientist to be involved
- Can integrate VSAM and IMS data through virtualization (DVM)
- Leverages z stack AI functions (zIIPs, Telum™, …)

An example:

Here is the account ID of someone who engaged in fraudulent activity.

Show me the 100 accounts most similar to this one

SQLDI Pro Differentiators

Unified AI for Structured + Unstructured Data

- Competitors focus on text chunks only.
- Orion uniquely merges relational attributes + free-text fields into a single semantic model.
- One embedding = complete enterprise view.

In-Place on IBM Z

- No need to move data to external vector stores.
- AI runs where the data already lives, preserving compliance, security, and performance.

Incremental Training

- Updates models using only changed data.
- Faster iteration, lower cost, and always up to date.

Built-In Model Explainability

- Critical for trust, audit, and regulated industries.

Hardware Acceleration

- Optimized with Telum II and Spyre for embedding and inference speed.
- Scalability ready for GenAI workloads.

Seamless Integration with IBM Ecosystem

- Works with IDAA for cost-efficient training and inference.
- Native connection to watsonx Assistant for Z for conversational AI (Future).

DVM Overview

Compatible Data Sources

- DB2 for z/OS
- SMF Records
- Db2 Analytics Accelerator for z/OS
- Sequential
- IBM IMS
- ADABAS
- VSAM
- z/OS
- ...and more!

Mainframe

Non-Mainframe

- ORACLE
- PostgreSQL
- APACHE hadoop
- teradata.
- Microsoft SQL Server
- Informix SOFTWARE

IBM Z

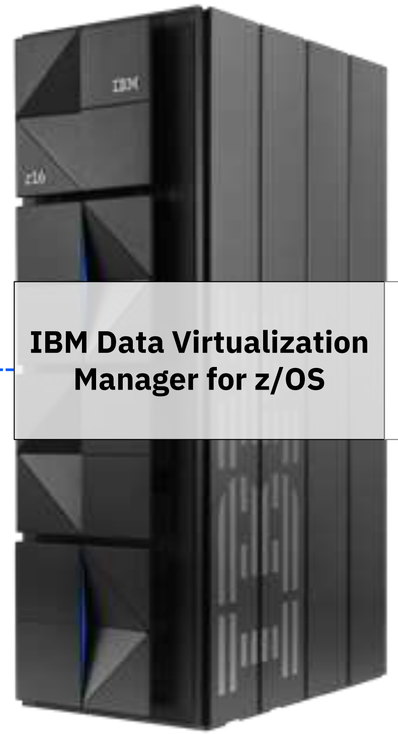

IBM Data Virtualization Manager for z/OS

JDBC
ODBC
REST

Data Virtualization

Data Targets

- Modernized Mainframe Apps
- Modern Data Platforms
  - Cloud Pak for Data
  - watsonx.data
  - presto
  - Spark
  - ICEBERG
  - Governance and Metadata
- Cloud/Mobile Apps
- BI Dashboards

IBM Z Data Integration Summary

- Direct access and virtualization
- Data synchronization / copy

DVM : Data Virtualization Manager for z/OS

LinuxONE

IBM Z

IBM Z Data
- Db2z
- IMS
- VSAM

Db2 and/or DVM JDBC

Db2z / IMS / VSAM

Db2 and/or DVM connector

Data Gate

Db2z data *only*

Db2 Warehouse

Db2 connector

Db2 OLTP

Db2 connector

Data Replication

kafka

Data Gate for watsonx

Db2z / IMS / VSAM

ICEBERG

Open Data Format on Object Store

Data Gate for Confluent*
Announcement: May 5th, 2026
GA: June 12th, 2026

CONFLUENT

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.

Data Gate – Transactional Use Case

- **Development of new cloud-native applications that requires Db2 for z/OS data, e.g., mobile banking applications**
  - Minimal operational impact on Db2 for z/OS
    - Working on cached copy
    - zIIP-eligible integrated synchronization workload
  - High performance requirements
    - Low-latency data synchronization protocol
  - *Transactionally consistent query results on latest data state*

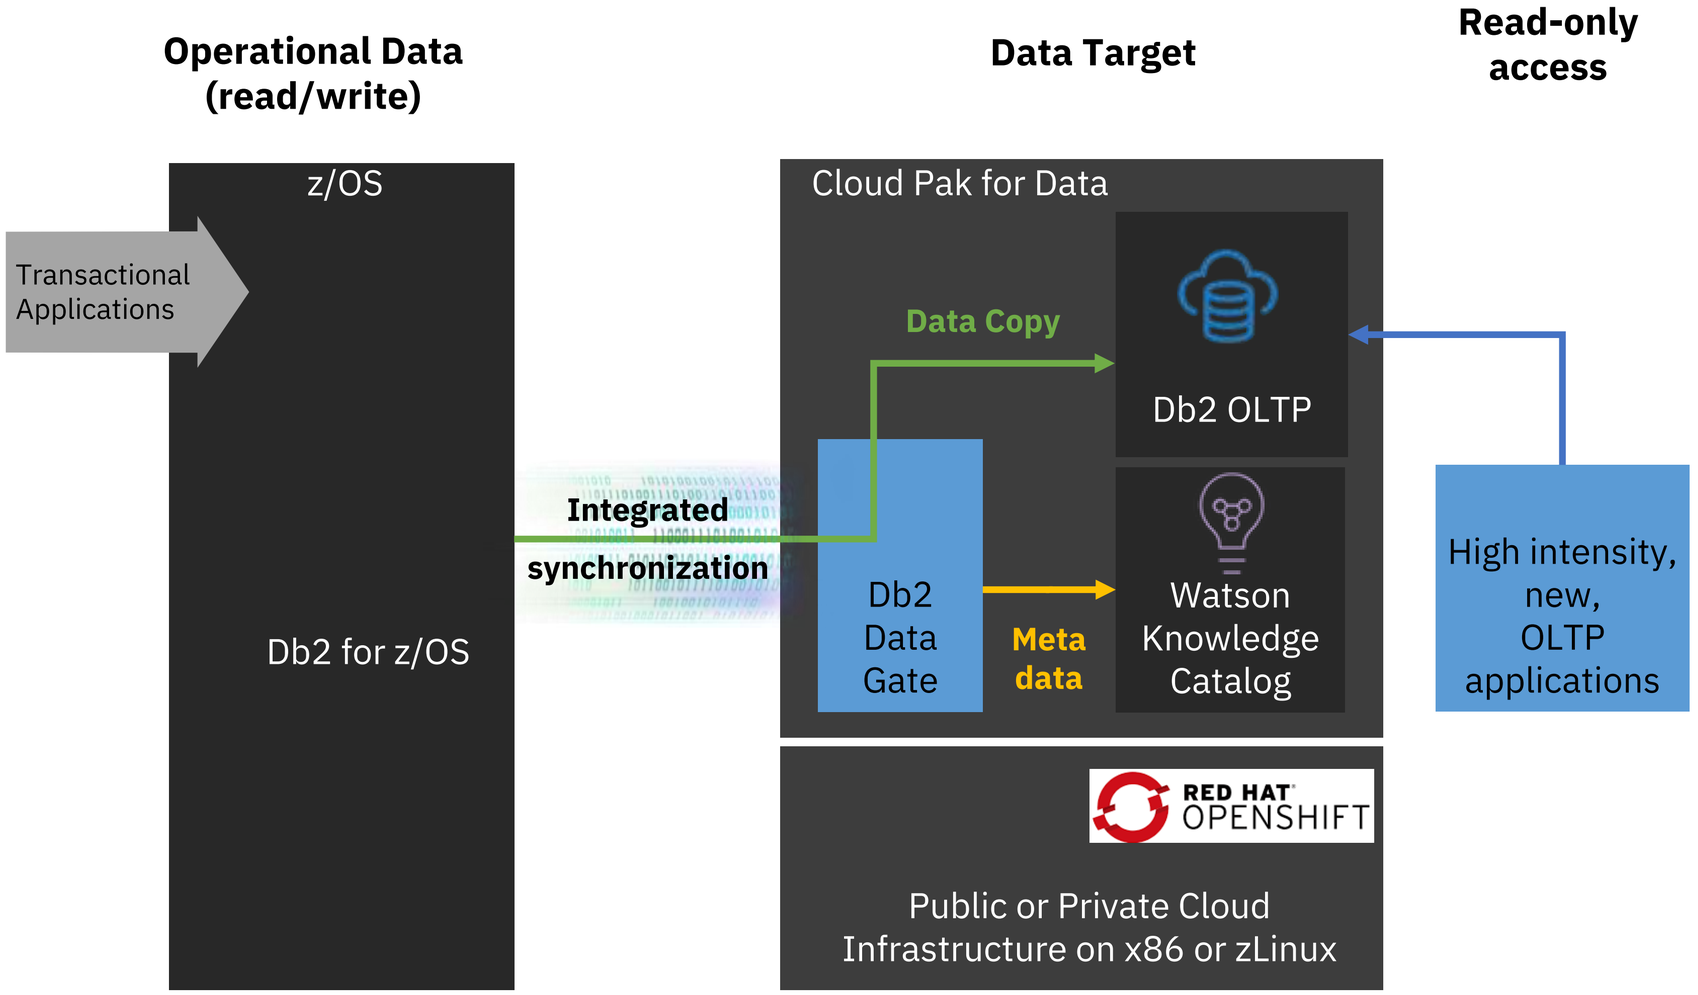

Data Gate – Transactional Use Case – Customer Example on IBM Cloud

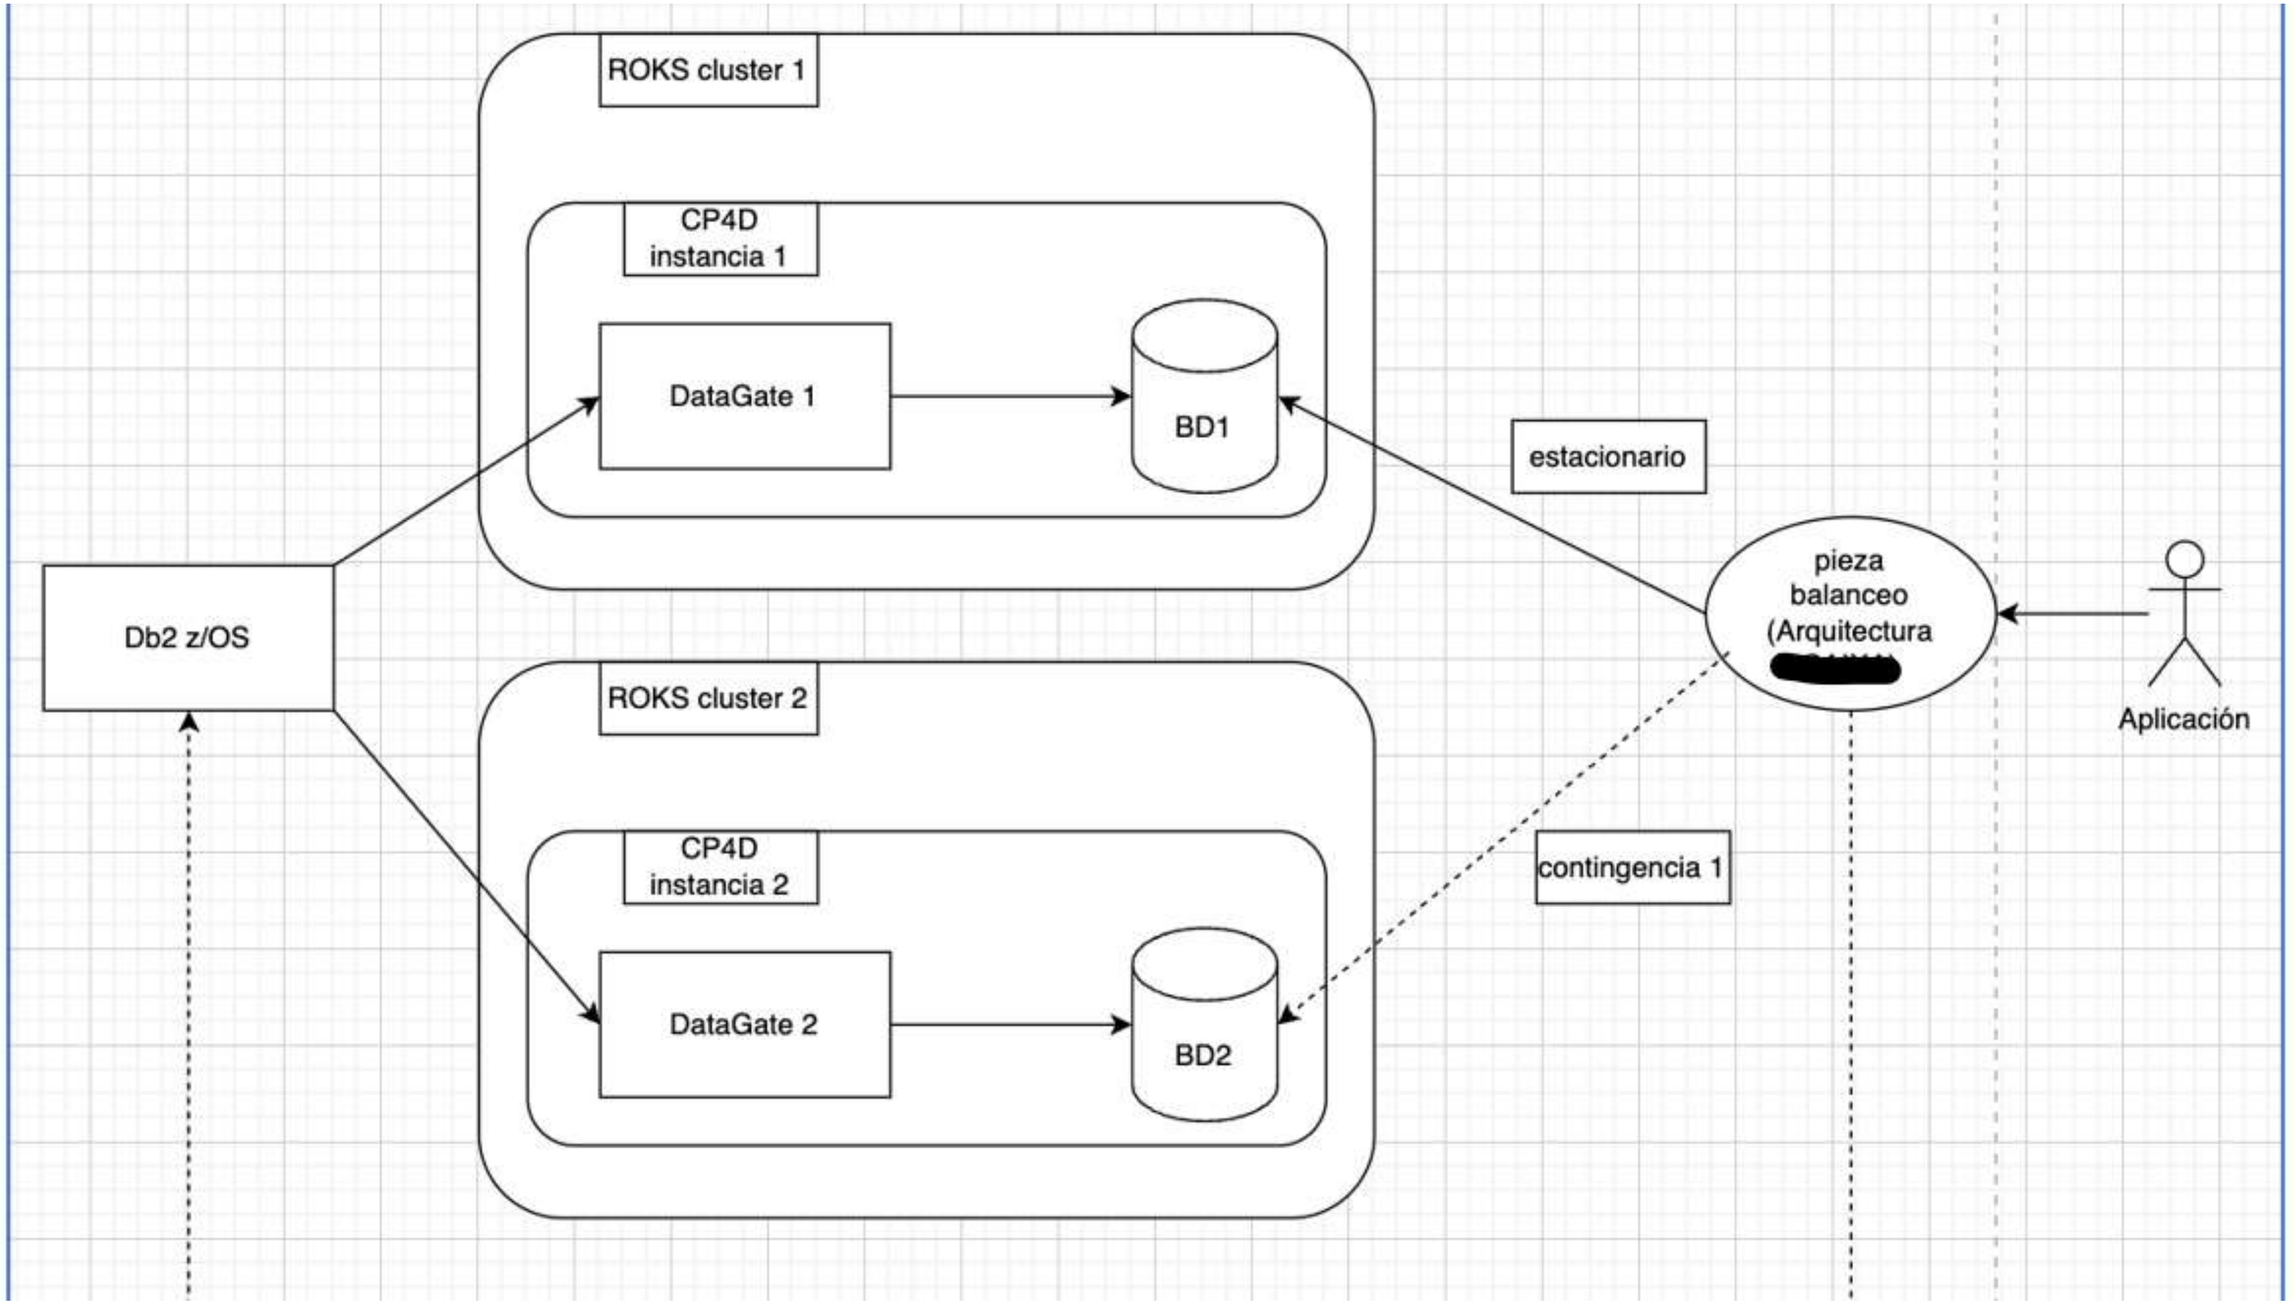

Data Gate – Transactional Use Case – Customer Example on Azure

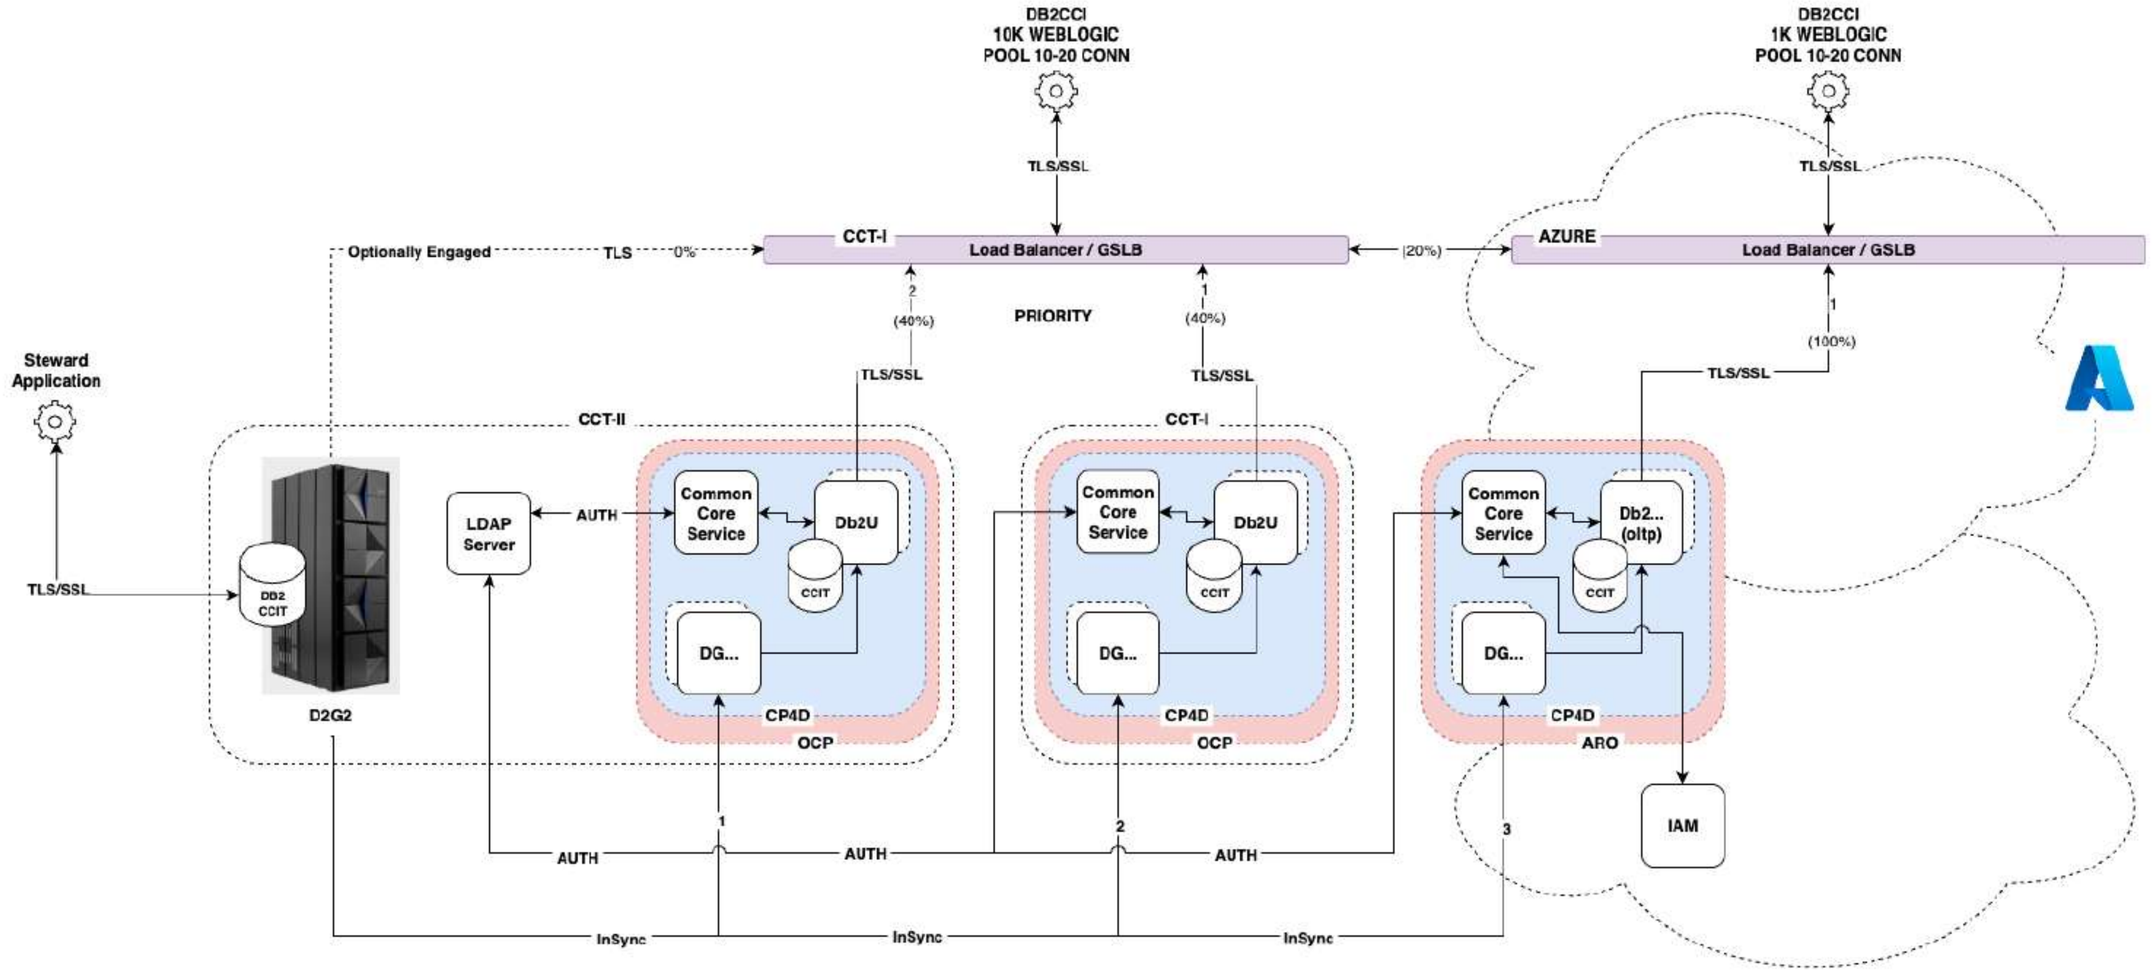

Data Gate – Transactional Use Case – Customer Example on x86 & Azure – Details

Business Case: Modernize access to Db2 for z/OS, ensuring scalability and resilience without compromising performance and governance. Cloud Native applications, such as application servers (WebLogic and etc), need to consume this data securely, with customized security certificate, in near real time, with support for high volumes of simultaneous connections and the ability to balance workloads across different sites on premises and the cloud.

Implementation: Data Gate Classic integrated with **Db2U (OLTP)** across multiple sites (CCT-I, CCT-II, and Azure/ARO), within **Cloud Pak for Data (CP4D)** running on **OpenShift** and **IBM Storage Fusion** (a.k.a. OpenShift Data Foundation - **ODF**).

Architecture: Symmetric production (PRD), test (HML), and development (DEV) environments. Load balancing workload is performed via Global Service Load Balancer - **GSLB with the proportion like (can be customized): CCT-I 40%, CCT-II 40% and Azure 20%**, with full failover to the cloud or on premises in contingency scenarios. If any unplanned issues occur, the final mile encompasses the connection between the application servers and the mainframe source, responsible for synchronizing data into Db2U.

Data Gate – Analytical Use Case

- **Development of new analytical cloud-native applications that requires Db2 for z/OS data**
- Db2 WH as target
  - In Memory, Columnar technology
  - IBM's premier analytics engine across many products

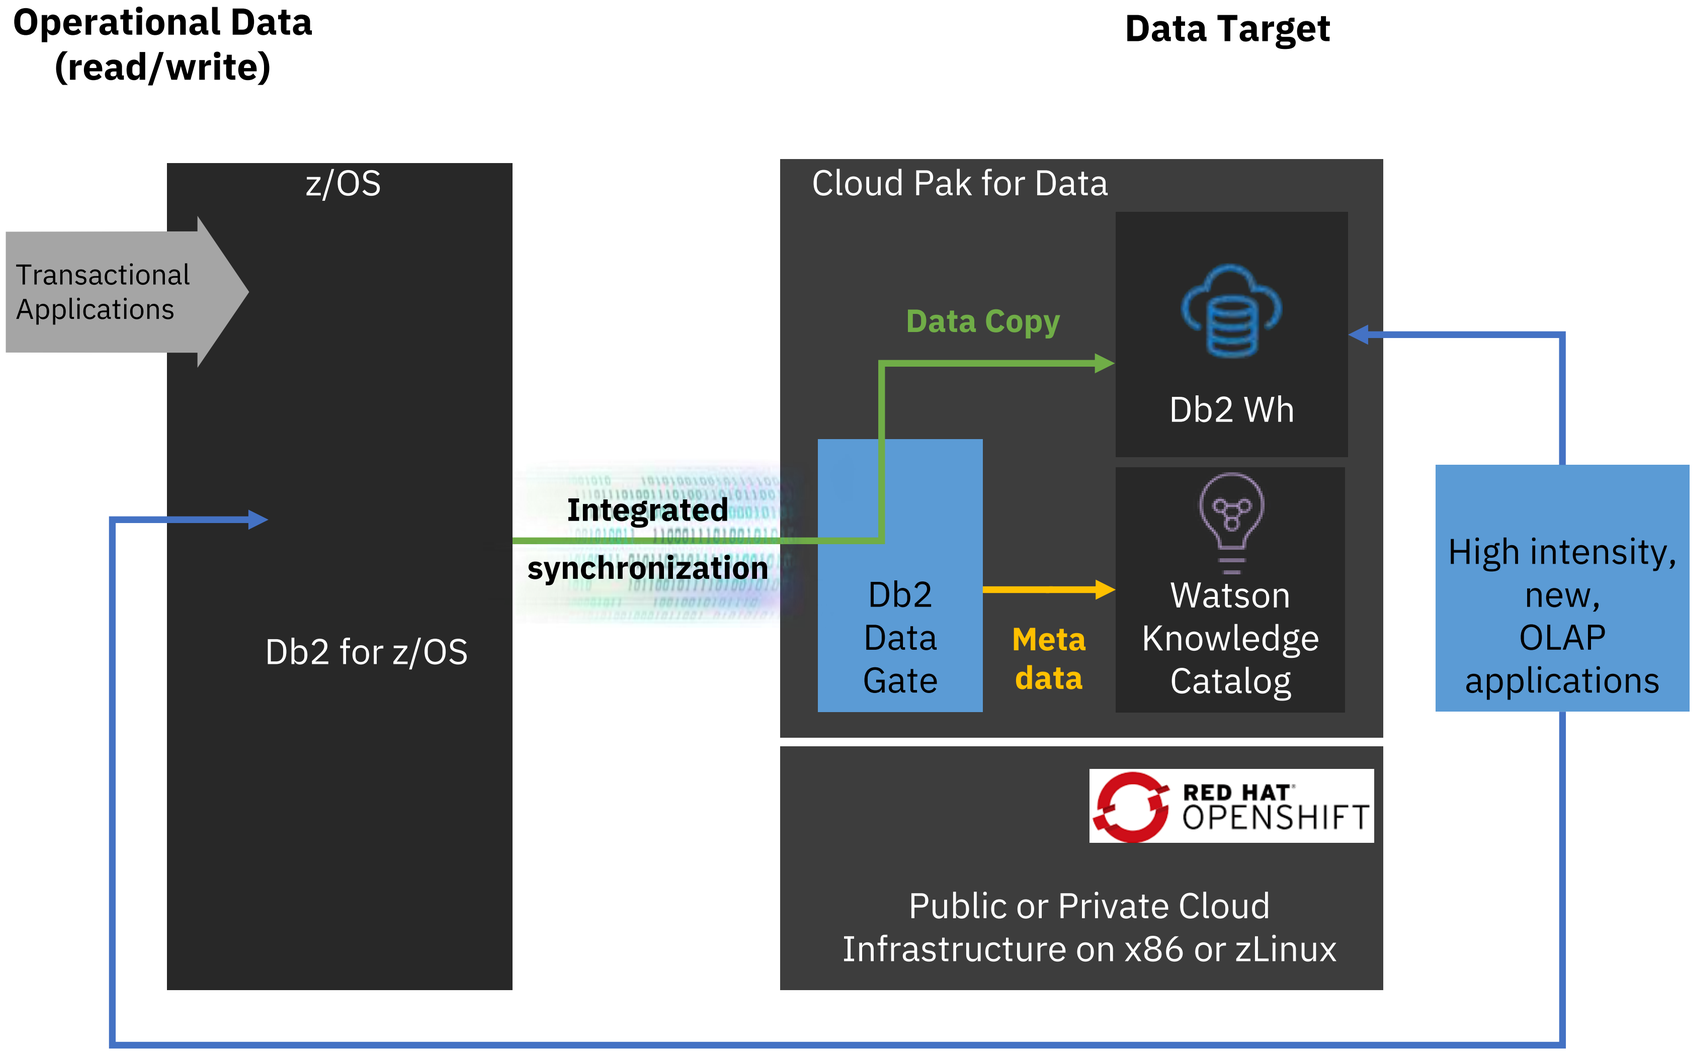

Data Gate – Combined use cases (API and Analytics)

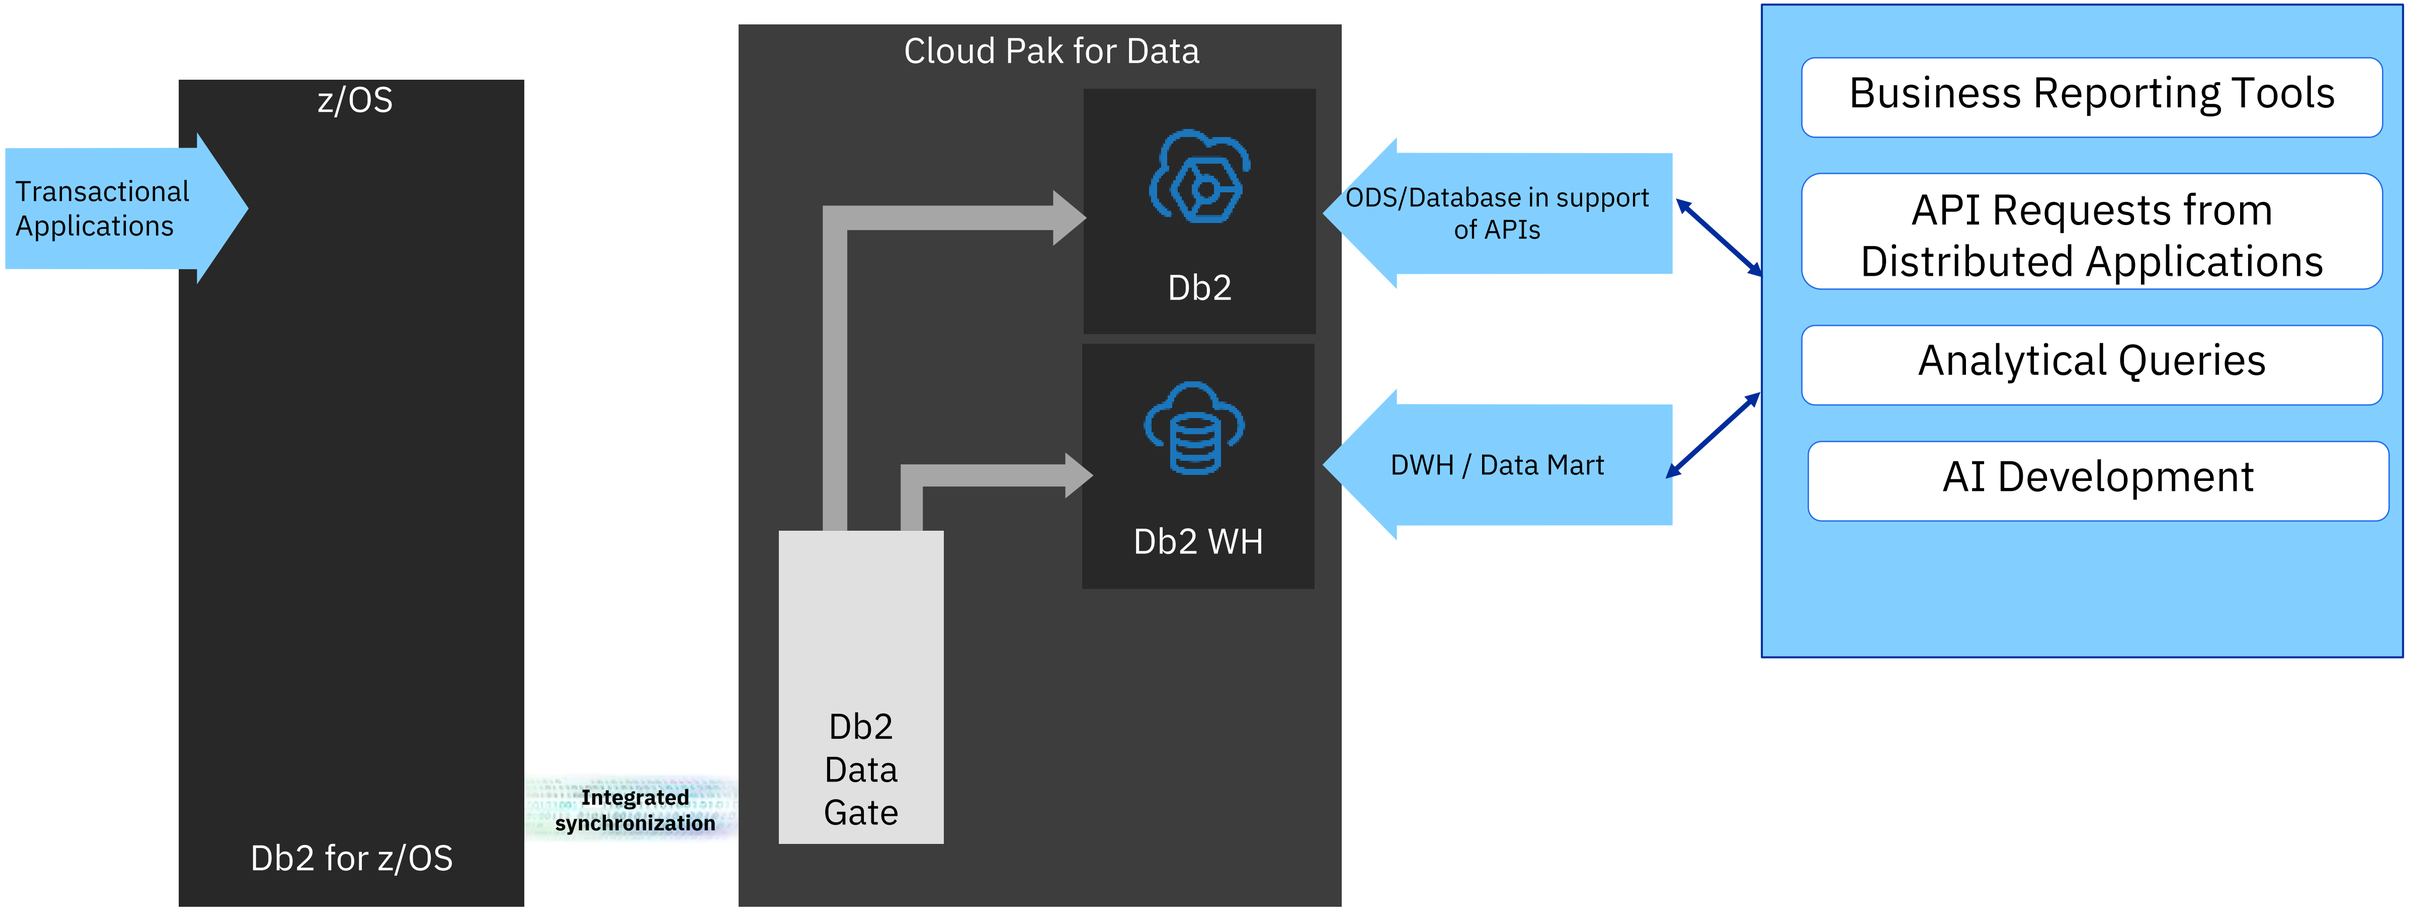

Replication Optimization (Part 1)

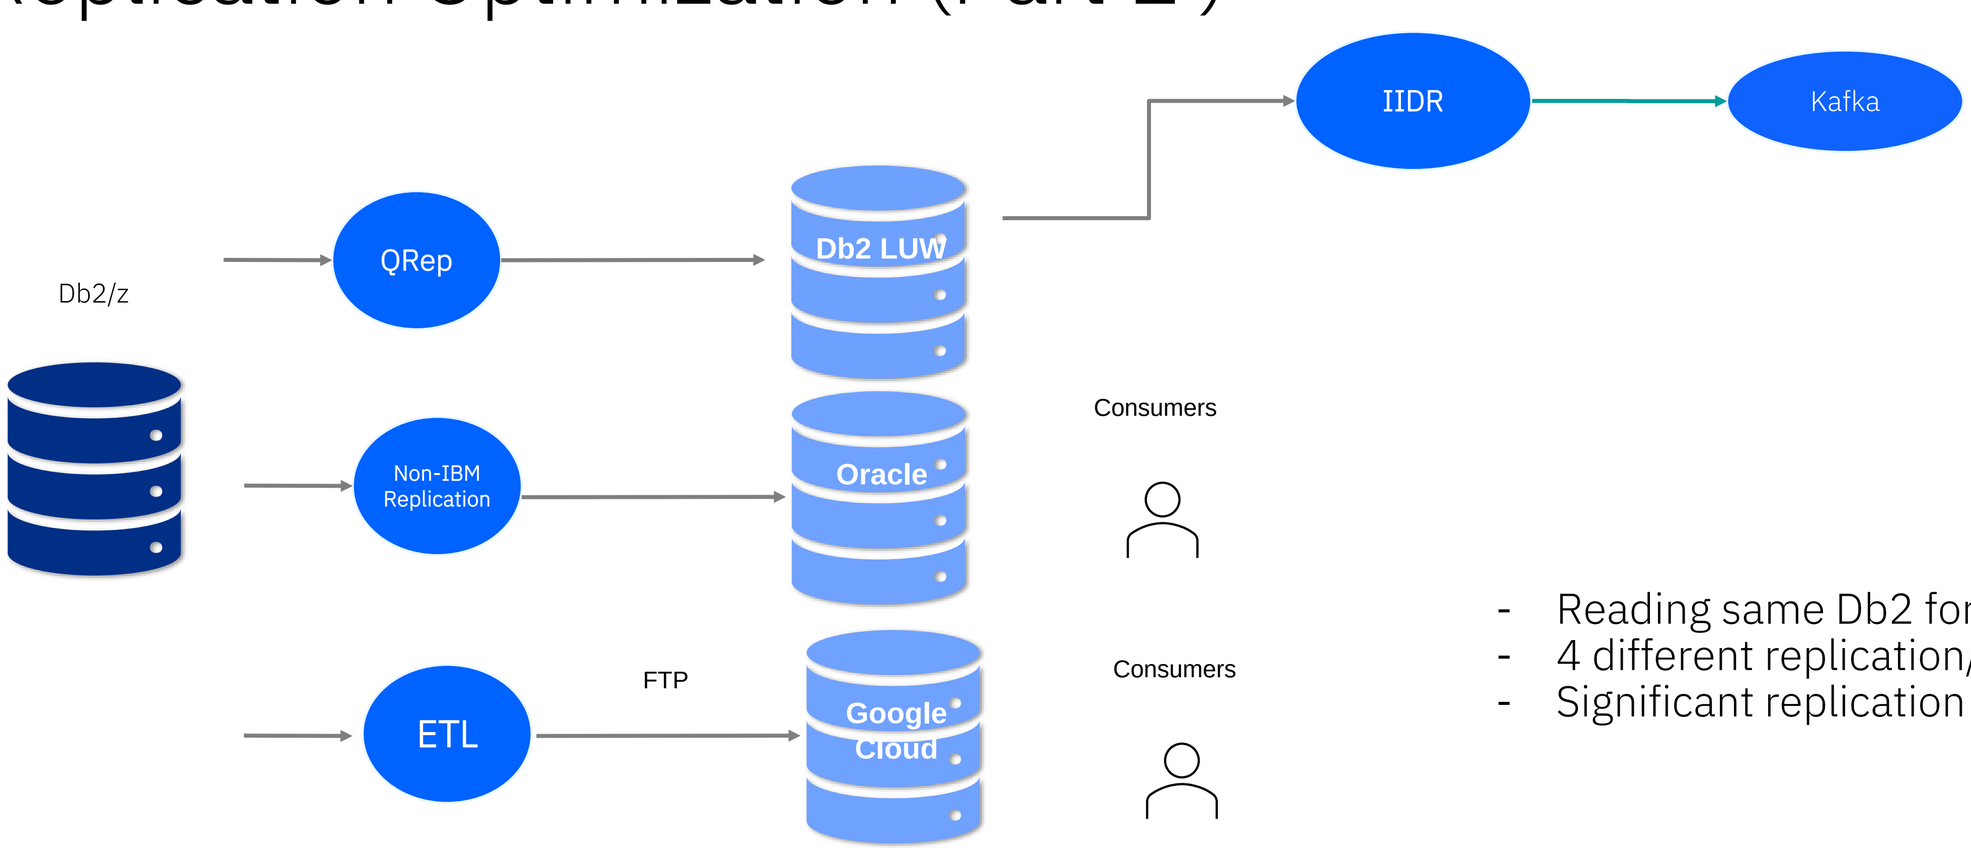

- Reading same Db2 for z/OS data 3 times
- 4 different replication/transformation tools
- Significant replication overhead to IBM Z

Replication Optimization (Part 2)

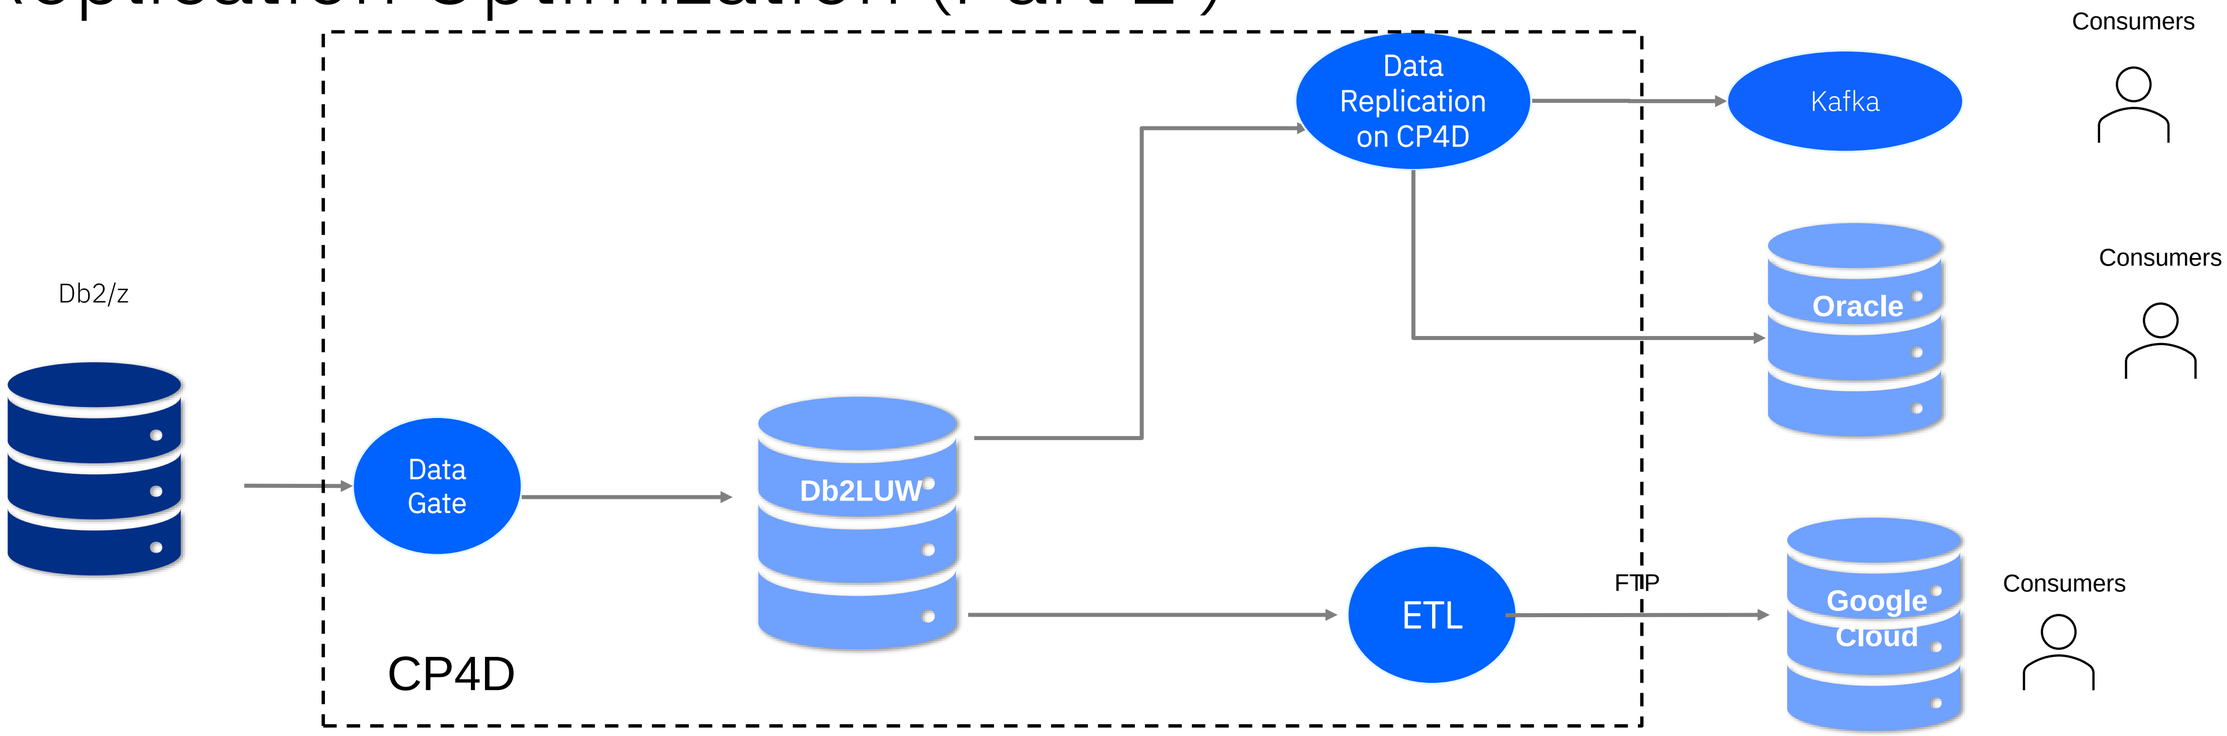

- Reading same Db2 for z/OS data once
- Eliminate one replication tool
- Minimum replication overhead to IBM Z
- Leverage zIIPs for replication on IBM Z

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.

Augmenting existing warehouse

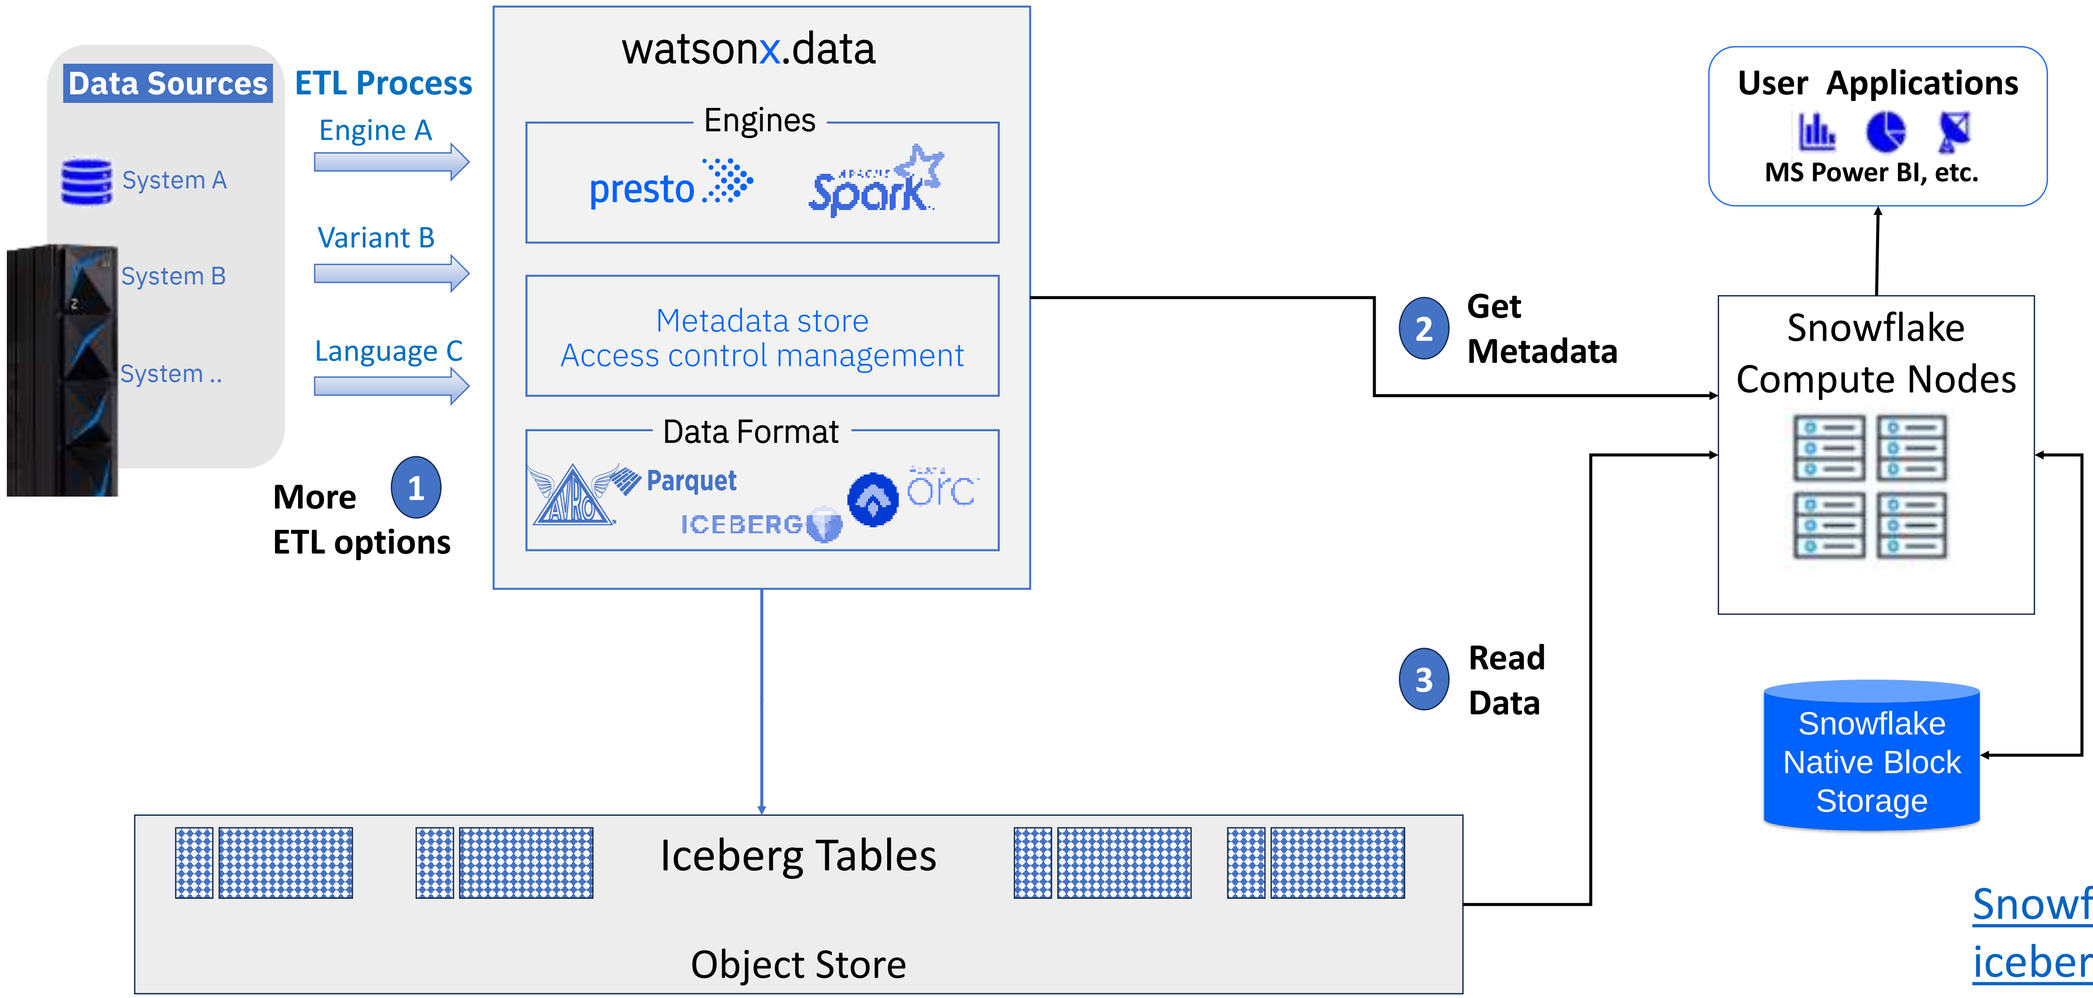

[Snowflake unified iceberg tables]

IBM Z and Salesforce Data Cloud Zero Copy Integration

IBM watsonx.data and Salesforce Data Cloud will allow businesses to access and utilize data from IBM Z mainframes and Db2 databases without moving or copying data

Description

Making Your Mainframe Data, Salesforce Smart and AI-Ready

- IBM Data Gate for watsonx integrates Z data into watsonx.data, making it accessible in Salesforce Data Cloud.
- The integration allows IBM and Salesforce customers to utilize data from both platforms for AI-driven insights and use in IBM's AI Agents and Agent Actions to unlock new opportunities for innovation..
- IBM will be the first Zero Copy partner to provide seamless data flow between IBM Z and the Salesforce Cloud

Value for customers

- Maximize mainframe data value: **C**onnect mission-critical Z data for CRM agentic use cases with enhanced trust, governance, cost-efficiency and customer experience
- Bring IBM Z data to modern data platforms like data lakehouse, enabling open table formats such as Apache Iceberg
- Streamline AI model training and analytics by exposing IBM Z data in a relational format, reducing overhead and enhancing throughput.

Key highlights

Expose IBM Z data to modern data platforms like data lakehouse

- *Async replication of Db2, IMS and VSAM through a single integrated solution enabling open table formats like Apache Iceberg*

Reduce MIPs on mainframe related to Db2 replication workloads vs. 3rd party tools

- *Replace legacy replication, warehouse and lake solutions with new pipelines that work natively with cost-effective, highly redundant object storage such as AWS S3, ADLS and IBM COS*

Streamline AI model training and analytics by exposing IBM Z data, including VSAM and IMS, in a relational format, reducing overhead and enhancing throughput.

High-level Architecture – Salesforce + IBM Zero Copy

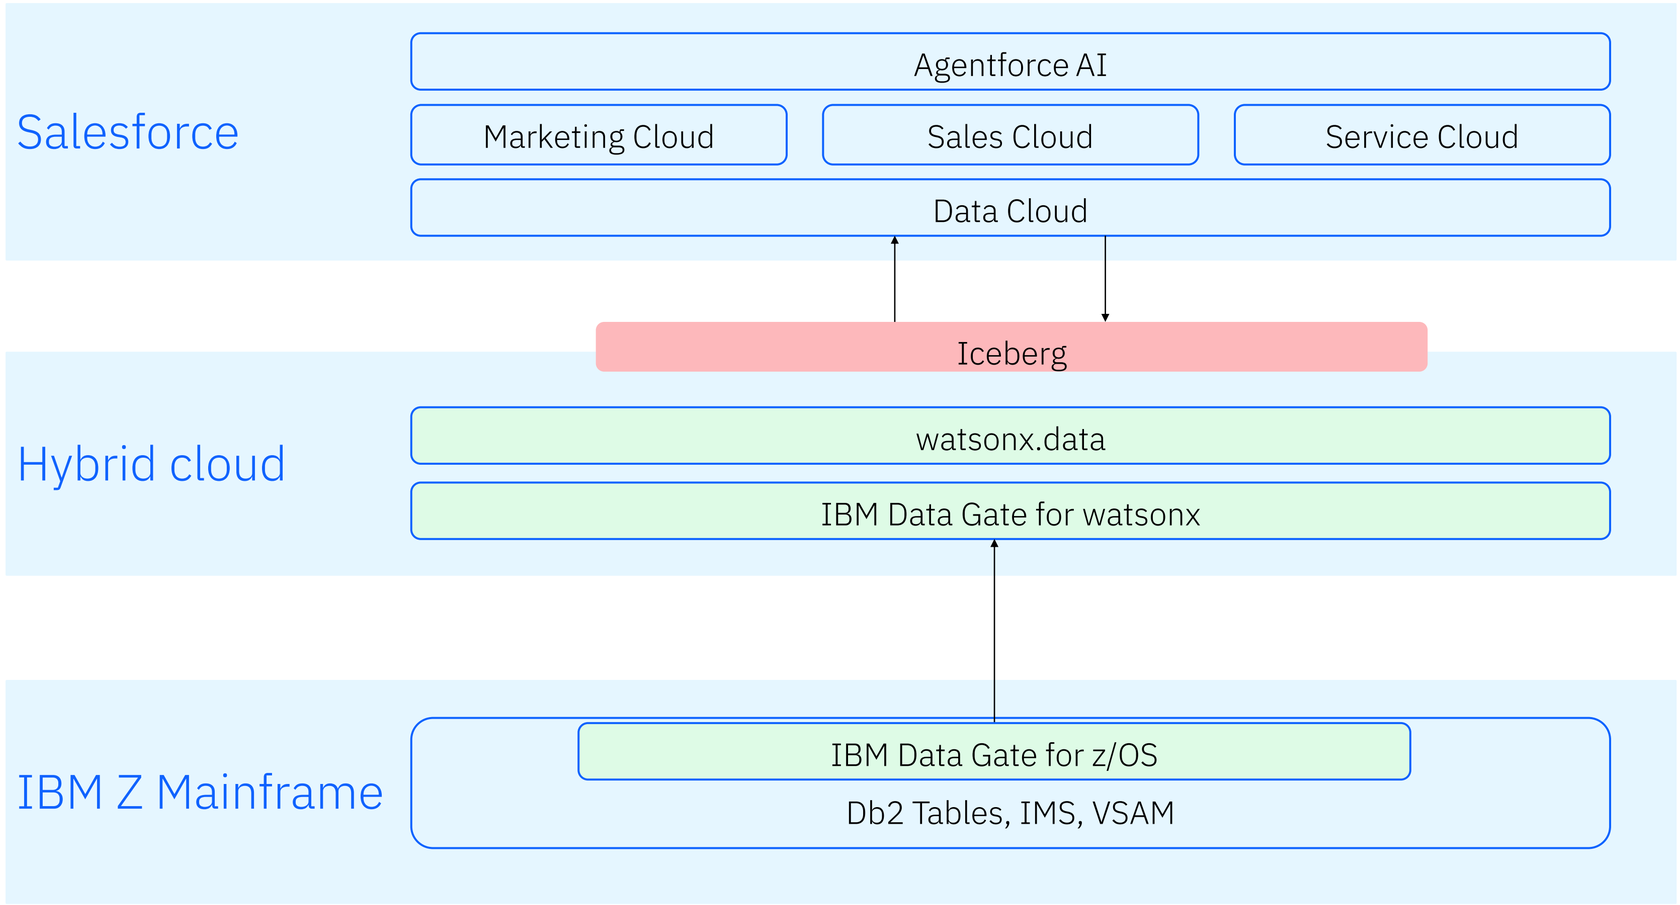

IBM Z Data Integration with Hyperscalers such as databricks

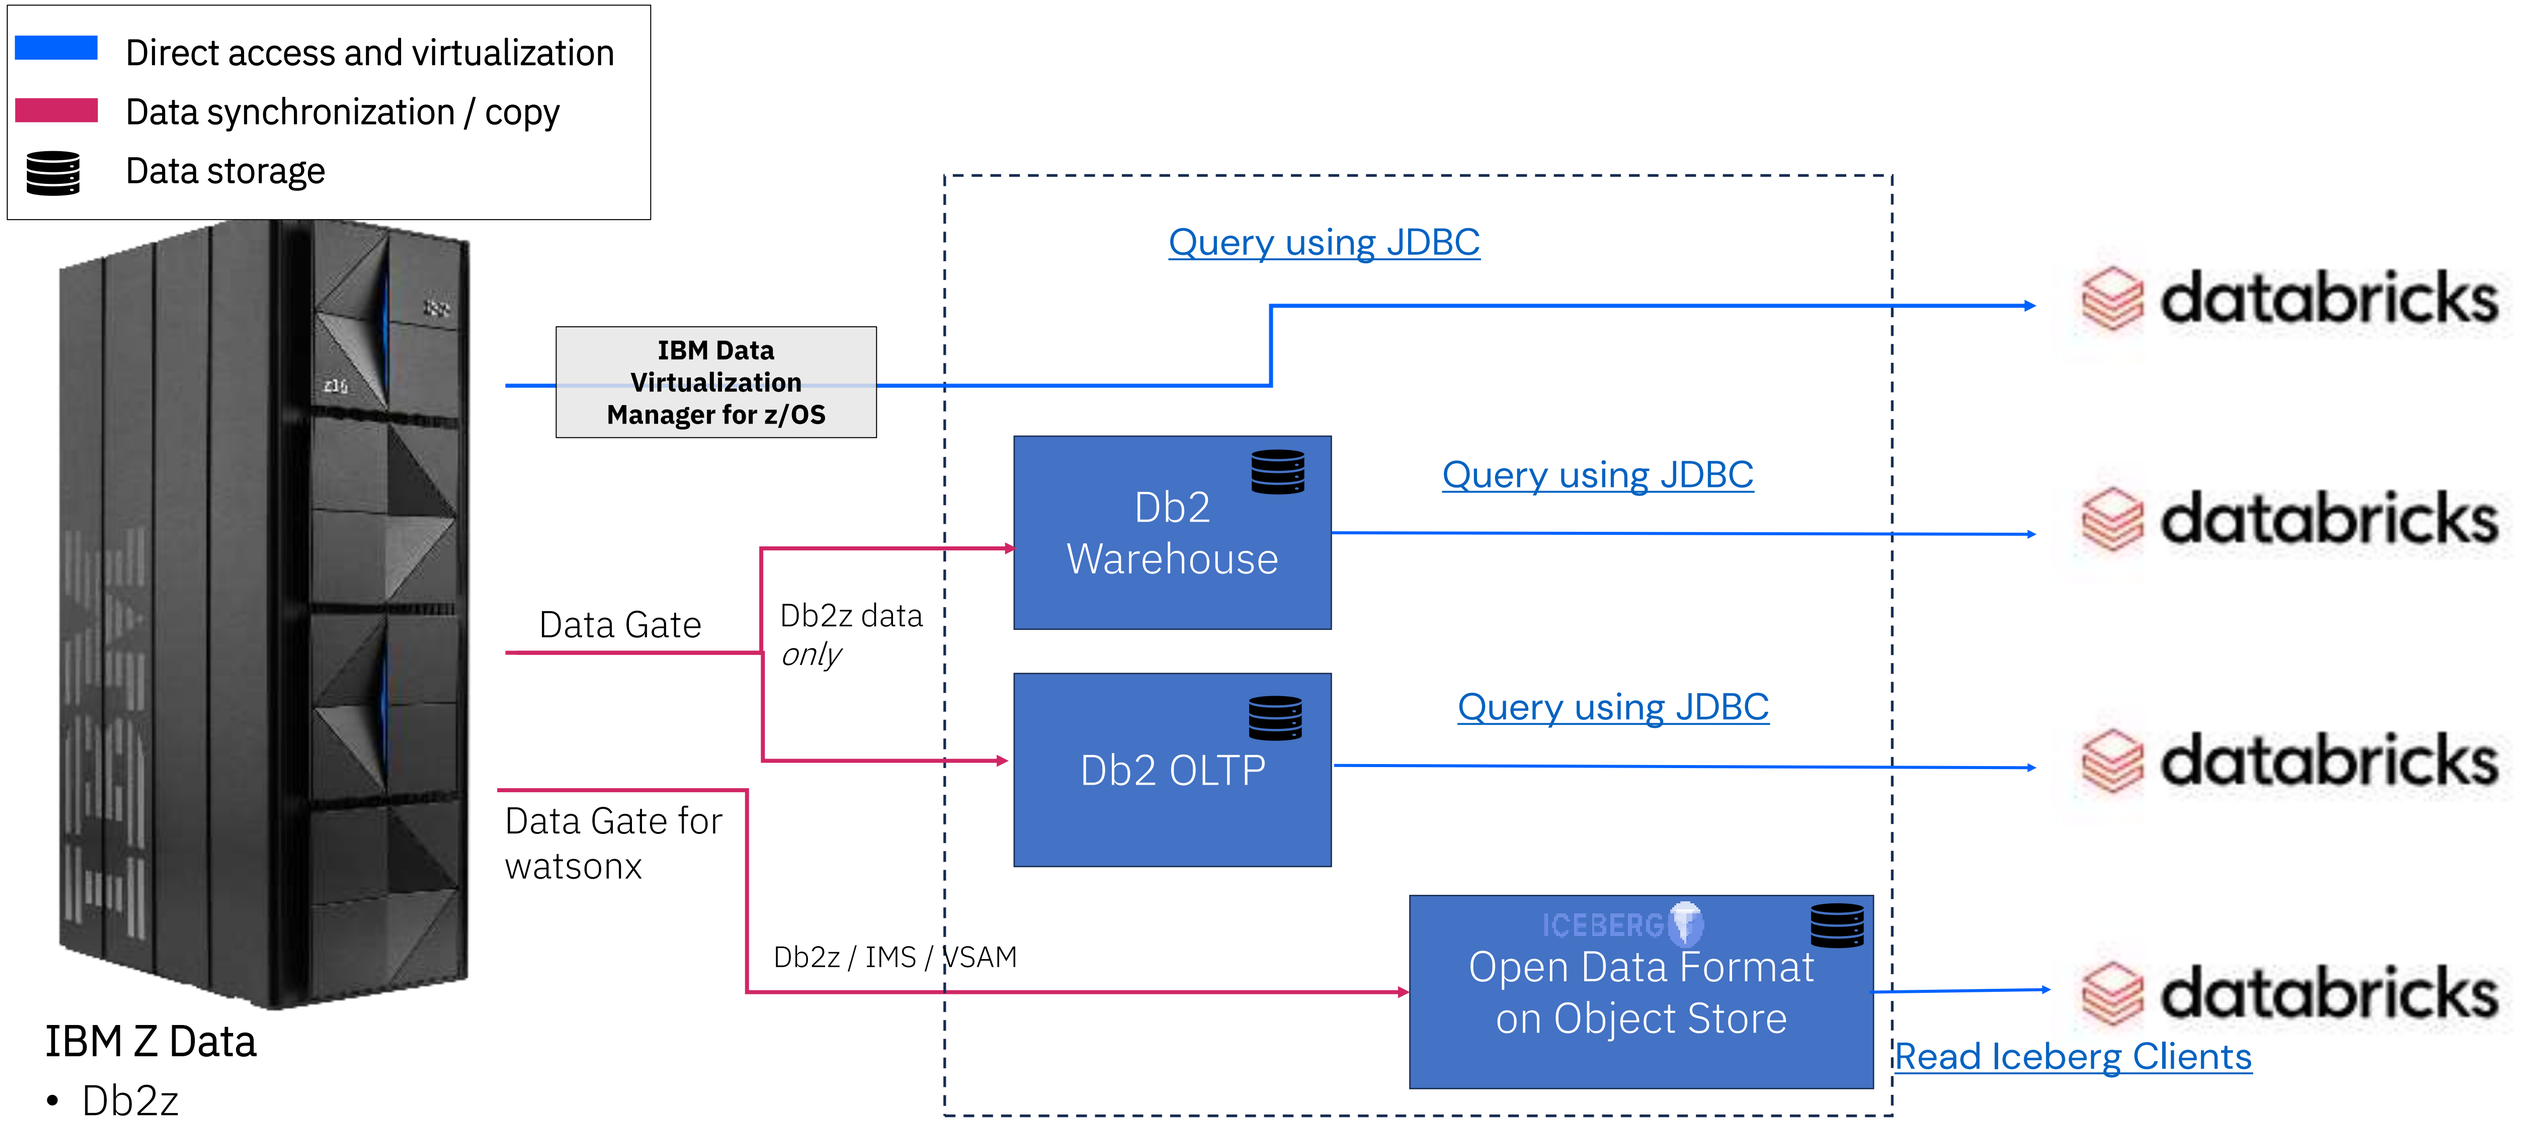

Kafka has become the de facto standard for data streaming

Created by the founders of Confluent while at LinkedIn

80%

Of the Fortune 100 estimated to be using Kafka

>100K

Organizations have adopted Kafka

>1,200

Kafka Improvement Proposals (KIPs)

>300K

Kafka meetup group members globally

Apache Kafka® is the de-facto standard for event-driven applications.

IBM Confluent is the complete, enterprise-grade data streaming platform built on Kafka

Confluent Platform for Z and LinuxONE

Integrate z/OS Applications with Confluent deployment on IBM Z Linux

Leverage high performance integration of **data sources eg. MongoDB, PostgreSQL with** Confluent on IBM Z Linux

Make Z Linux as Strategic enterprise event platform for enterprise wide applications on all platforms

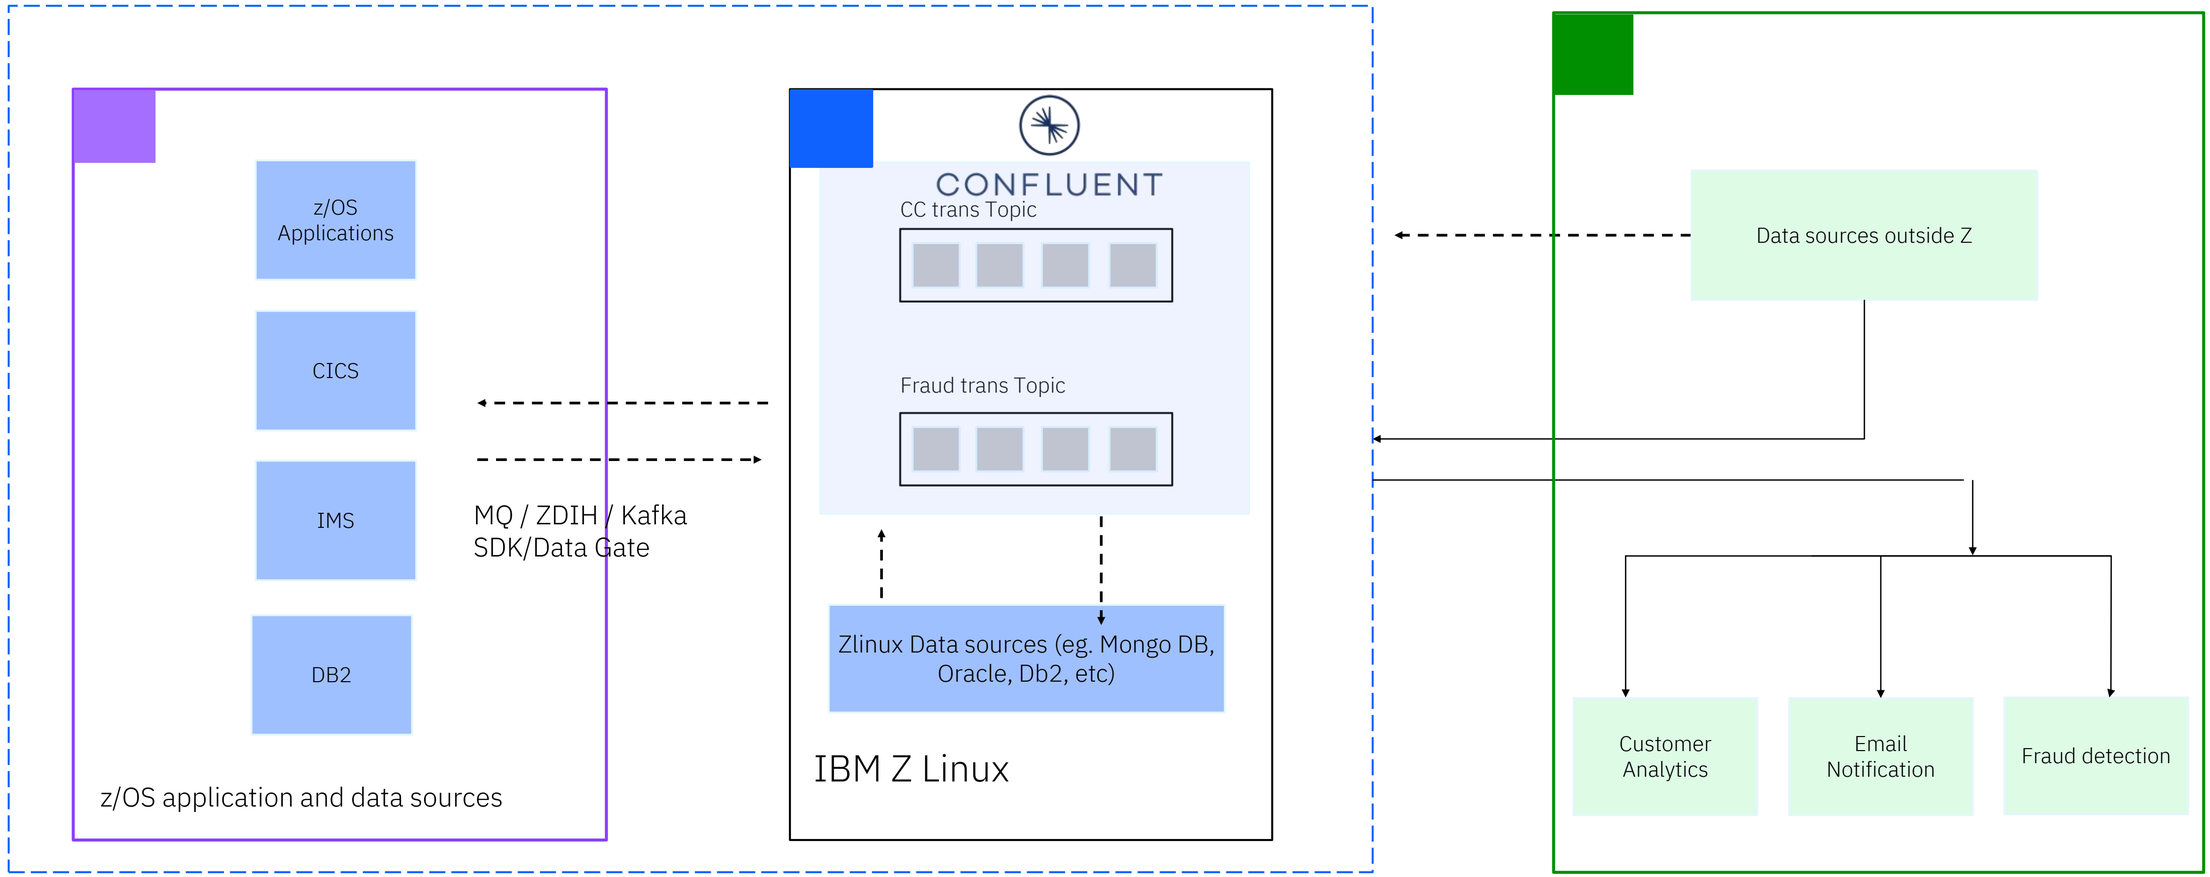

Data Gate for Confluent – High-level Overview

- **Consumption of Db2 for z/OS data in open data formats in Kafka environments**
  - Standardized integration into Kafka-based ecosystems
    - ➢ Interoperability for data consumers by data propagation in open Debezium event format
    - ➢ Source connector for Kafka Connect clusters using built-in Data Gate technology
    - ➢ Operations via Kafka Connect
  - Minimal operational impact on Db2 for z/OS
    - ➢ Working on cached copy
    - ➢ zIIP-eligible integrated synchronization workload
  - High performance requirements
    - ➢ Low-latency data synchronization protocol

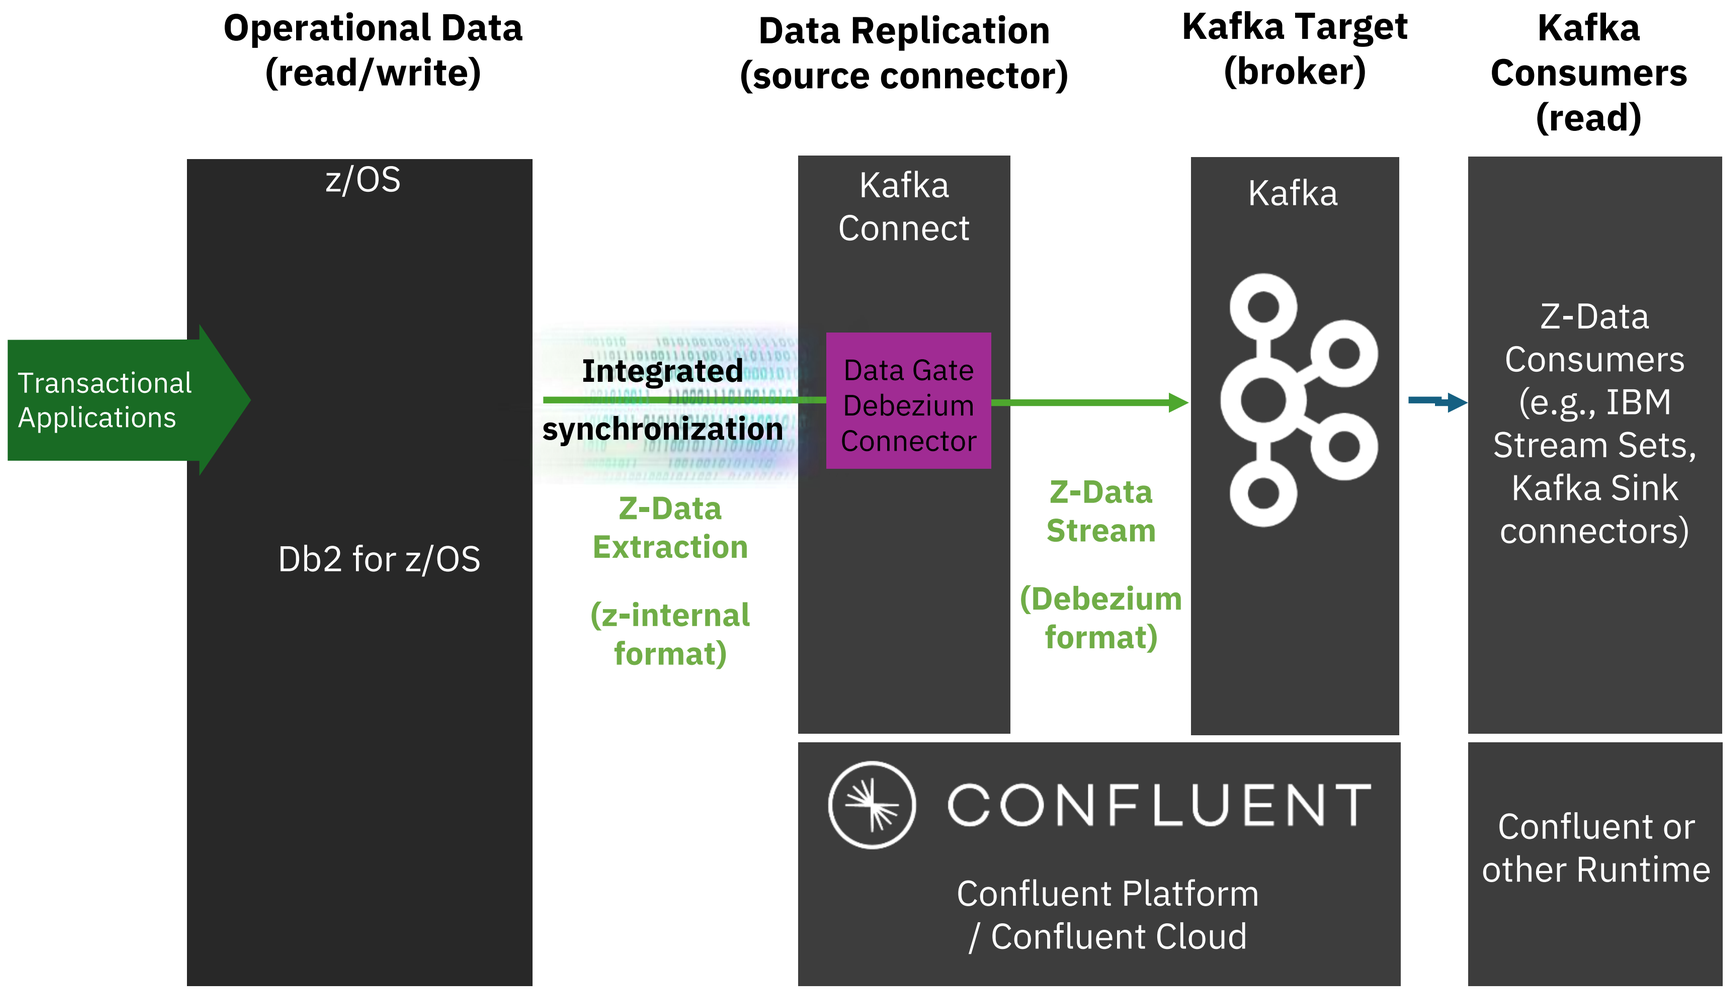

IBM Data Gate for z/OS – Version 3.2 with CPD version 5.3

Remote Db2 Support

Classic Data Gate *without* Remote Db2 Support

- Minimal operational impact on Db2 for z/OS
  - Working on cached copy
  - zIIP-eligible integrated synchronization workload
- High performance requirements
  - Low-latency data synchronization protocol
- Transactionally consistent query results on latest data state
- Data Gate propagates Db2 for z/OS data to a target database on Cloud Pak for Data
- Data Gate and Db2 must be in the same CP4D cluster

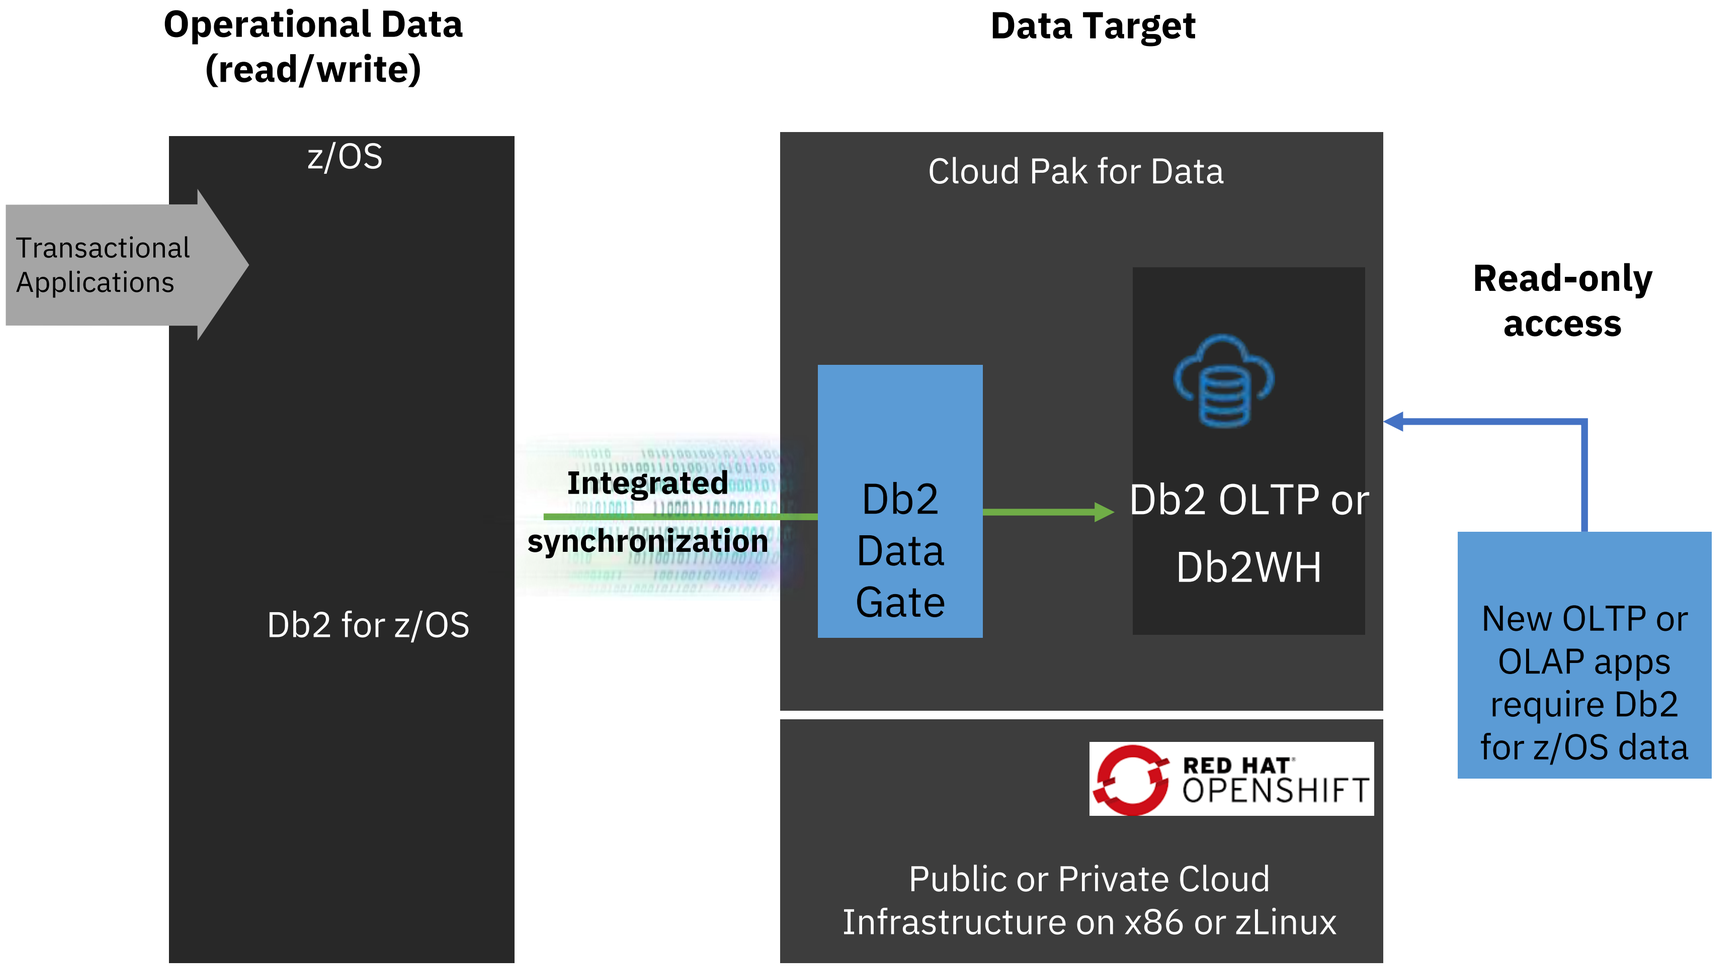

Classic Data Gate *with* Remote Db2 Support

- Db2 instance, may reside in a different Cloud Pak for Data cluster or on a remote external Db2 instance.
- The target Db2 instance is configured manually before connecting to the Data Gate instance which is outside of Data Gate scope.
- There is no plan for configuration UI, just provide the instructions via the documentation
- Support for Db2 OLTP and OLAP target database instances since 5.3.1
  - Query acceleration and archiving use cases are targeted for 5.4 in June
- Db2 OLTP target database authentication with username/password only
- Db2 target database instance can only be used as a target of a single Data Gate instance. This is also currently investigated if the same Db2 instance can be used as the target for multiple Data Gate instances with isolation on the database name.
- Limited diagnostics and monitoring capabilities
- No IBM Knowledge Catalog integration

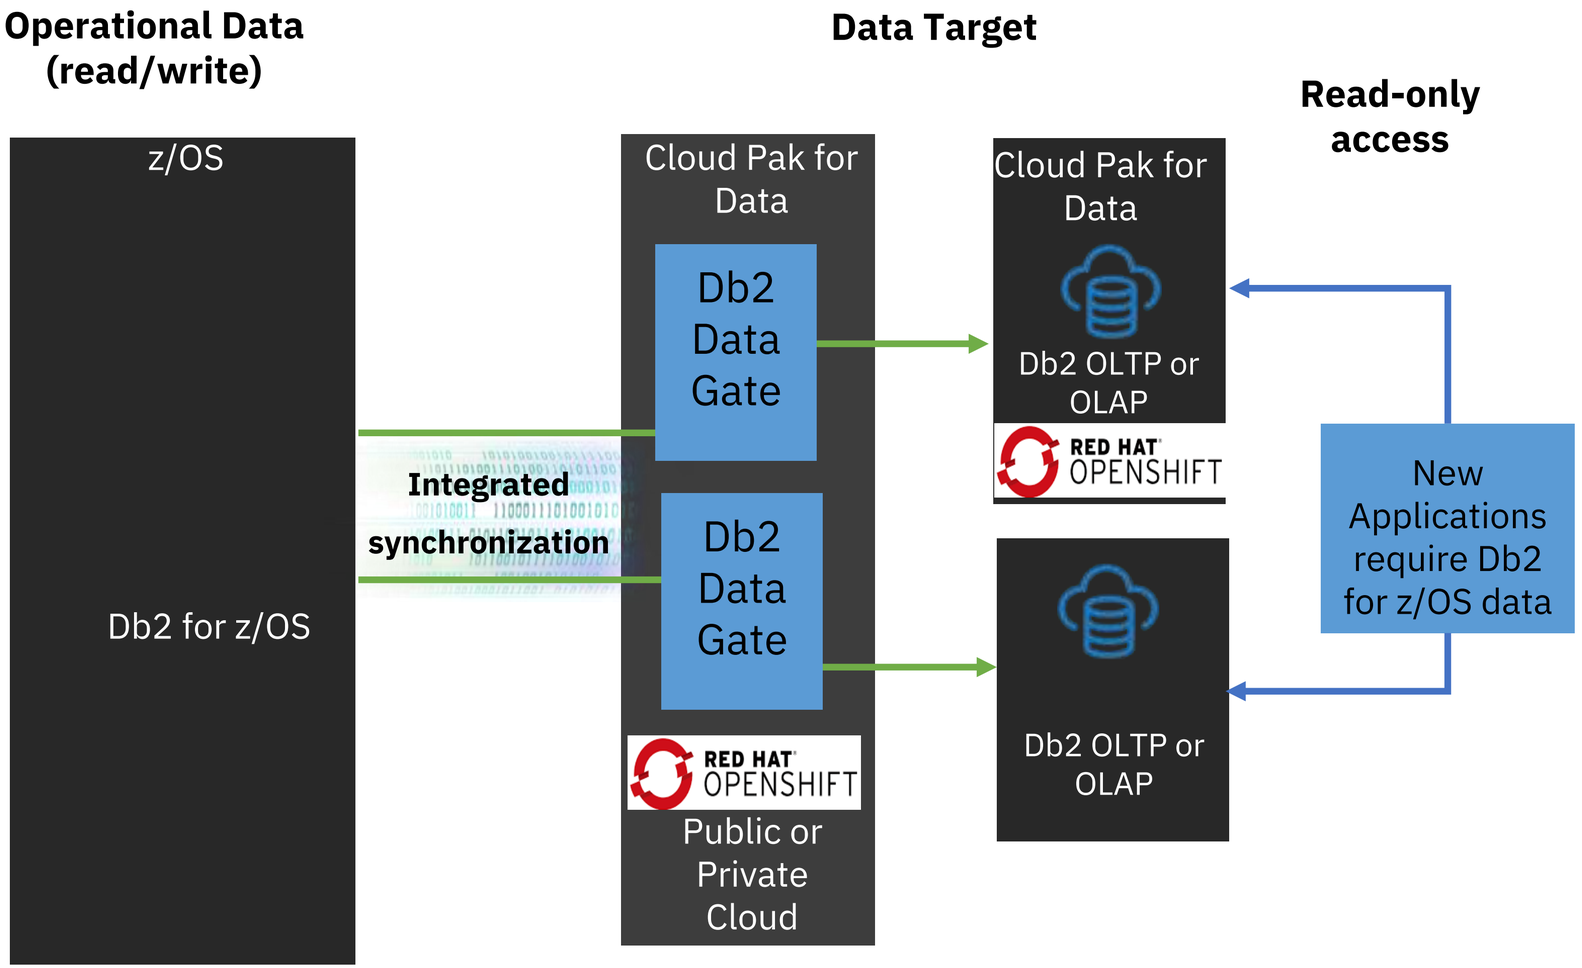

Architecture
- Db2 for z/OS source
- Data Gate as part of CP4D running on OpenShift cluster (x86 / s390x)
- Target Db2 OLTP or OLAP on any platform with network connectivity

[Connecting to a remote Db2 instance - IBM Documentation]

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 (Orion)

- 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)
zDIH (for caching and decoupled patterns)

- 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