

# Landing Db2 for z/OS, IMS and VSAM data in Iceberg tables

aka

**Leverage your mainframe data for AI and analytics by landing Db2 for z/OS, IMS and VSAM data directly in Iceberg tables on object storage**

Rickard Linck  
Brand Technical Specialist  
Information Architecture  
Data & AI, IBM



# Please note

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.

The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

# Notices and disclaimers

© 2024 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.

**U.S. Government Users Restricted Rights – use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

Information in this presentation (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

**Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

It is the customer’s responsibility to ensure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM’s products. IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at: [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Leverage your **mainframe data** for **AI and analytics** by landing **Db2 for z/OS, IMS and VSAM** data **directly** in **Iceberg tables** on **object storage**

## Why this title?

Mapping the words from the client driver for a new Solution (as stated in a customer PoT)

- **Objective:** Enable seamless access to **IBM Z Db2 for z/OS and VSAM data** within a **data mesh** for advanced analytics purposes.
- **Analytics and AI:** Serve **analytics and AI applications** without increasing the load on the System of Record.
- **Data Storage:** Utilize the **Iceberg open-source table format** to store data on **object storage**, benefiting from its **scalability, durability, and cost-effectiveness**. This ensures the data is **easily consumable by multiple analytics tools.**
- **Additional Considerations** (for all possible solutions):
  - **Cost Implications:** Including expenses related to making data available, particularly considering IBM Z gCPU usage.
  - **Licensing Requirements:** Necessary licenses for all tools and platforms involved.
  - **Human Resources:** Evaluate the effort required from personnel to implement and maintain the solution.
  - **Intermediary Components:** Consider tools and **staging storage solutions** that will act as intermediaries in the data flow.

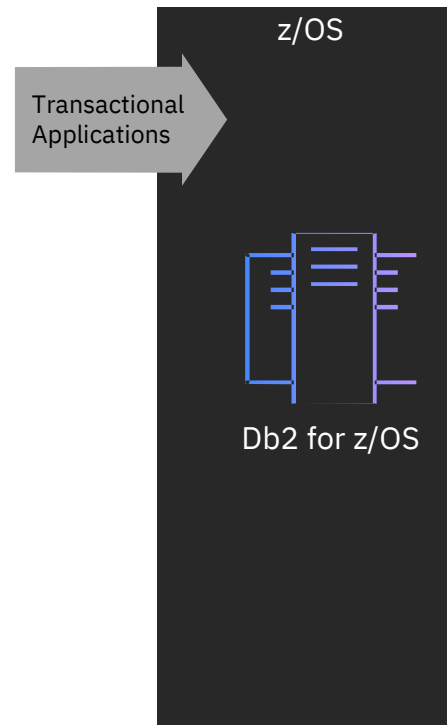
Let us revisit Data Gate

# Data Gate core values

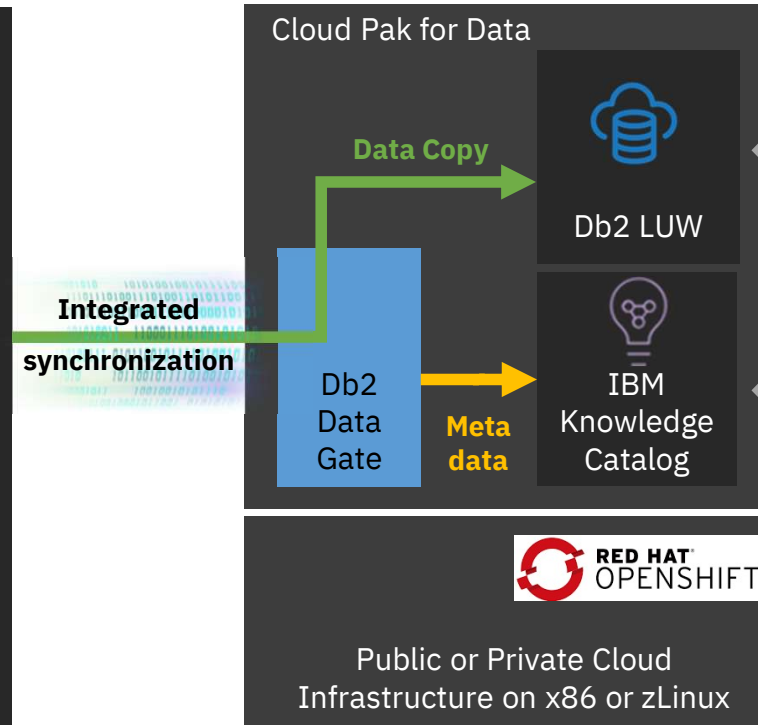
## Mission: z data modernization

- High-performance replication with Integrated Synchronization
- Low footprint on Z and zIIP eligibility
- Enable transactionally consistent queries on a synchronized cloud
- Containerized cloud-native application
- Metadata publishing
- Db2 for z/OS data only

## Operational Data (read/write)




## Data Synchronization (unidirectional)



## Data Target

## Read-only access



# Db2z: Initial Load and Integrated Synchronization performance

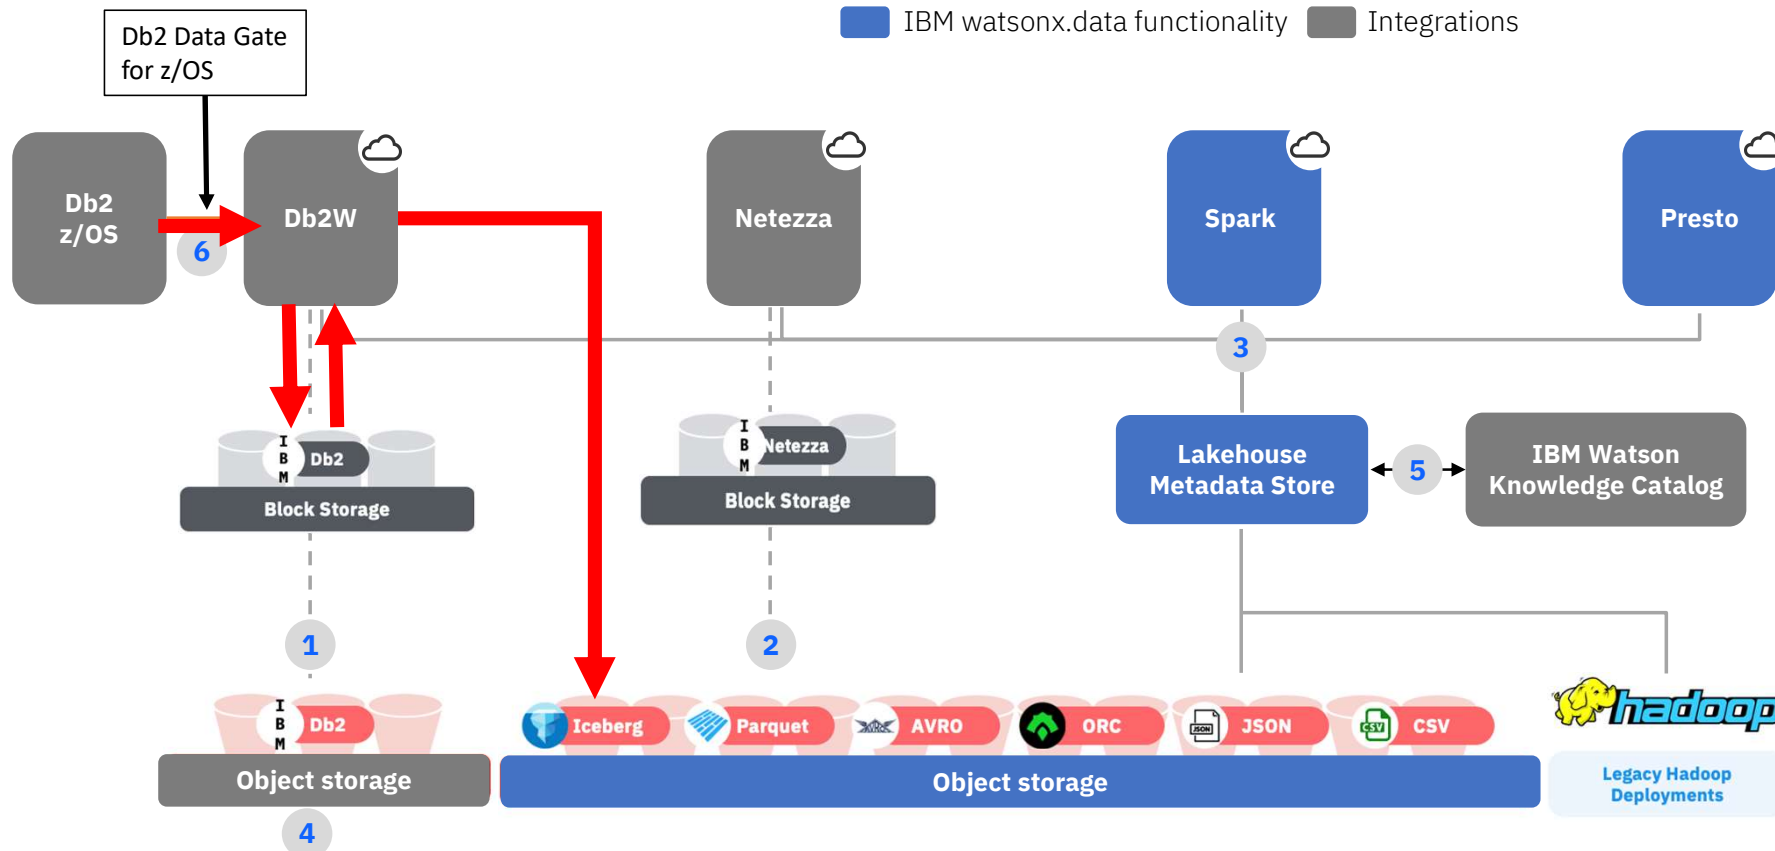
## Initial load

- **Parallel unload** of table partitions on the Db2 for z/OS side to maximize throughput
- **Parallel insertion** to target Db2 LUW tables with NLI(OLTP) and Reduced Logging(Warehouse)
- Transfer to Data Gate in Db2 for z/OS **internal format**
- High performance **code page conversion** to Unicode

## Synchronization

Workload	Transactional (Number of Rows, - avg. latency, max latency)	Analytical (Number of Rows - avg. latency, max latency)
INSERT	200K - 1.5s, 2.8s	200K - 1.3s, 1.8s
UPDATE	175K - 2.9s, 4.4s	150K - 2.1s, 2.6s
DELETE	200K - 1.2s, 2.1s	200K - 0.8s, 1.6s
MIX IUD	200K - 2.1s, 4.4s	200K - 1.5s, 2.5s

# Pre-2024: The integrated IBM watsonx.data ecosystem for maximum workload coverage and optimal price-performance



- 1 Warehouses can access data in the lakehouse
- 2 Easily Promote data between the warehouse and lakehouse
- 3 Query routing service, multiple engines can access same data lake data
- 4 The lakehouse can access data residing in Db2/Netezza
- 5 WKC policies enforced by the lakehouse via metadata service
- 6 Analyze Z data easily and securely with Db2 for z/OS Data Gate

At this time, Iceberg writing was too slow to keep up with mainframe transactions so Db2 for z/OS data lands in Db2 (LUW) instead and the data could then be “promoted” to Iceberg tables



# Data Gate variants in 2024


- **Data Gate on Cloud Pak for Data**

- Db2 OLTP for high-volume transaction workloads
- Db2 Wh for analytical applications in Cloud Pak for Data

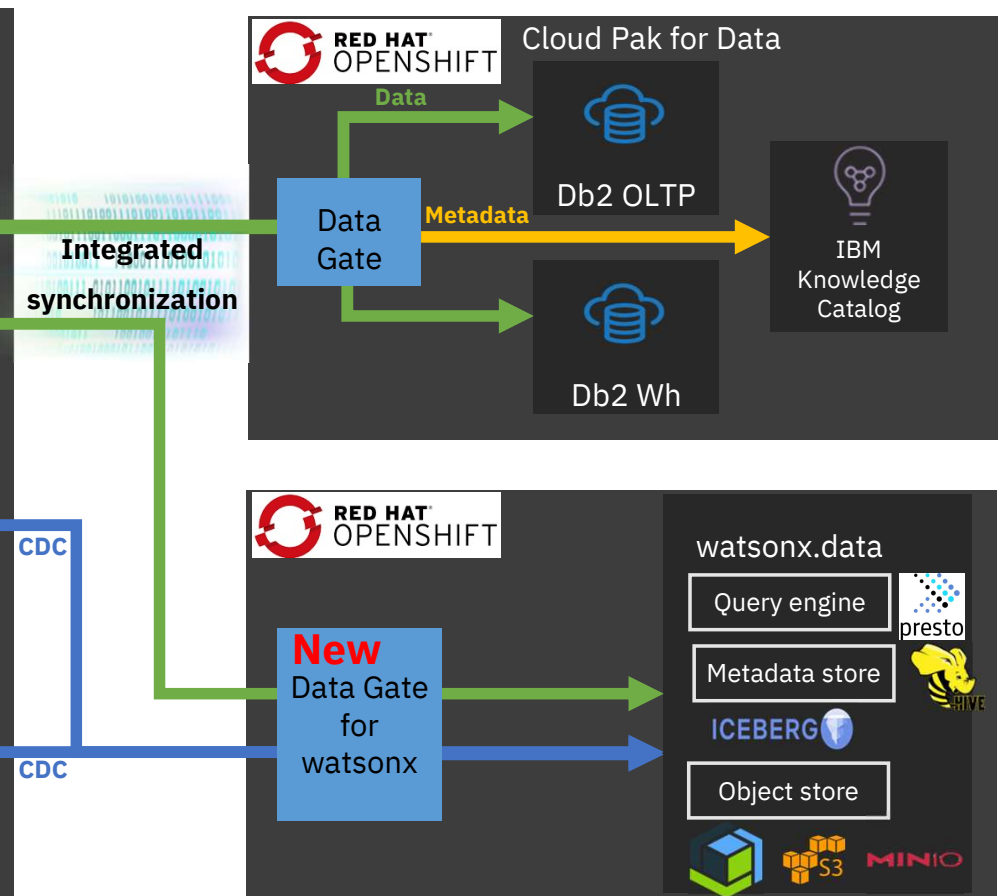
- **New Data Gate for watsonx**

- For analytics and ML in watsonx.data
- Combination with Spark streaming for streaming queries or writing to further data stores
- Db2 for z/OS, **IMS**, **VSAM**
- Target format is **Apache Iceberg**

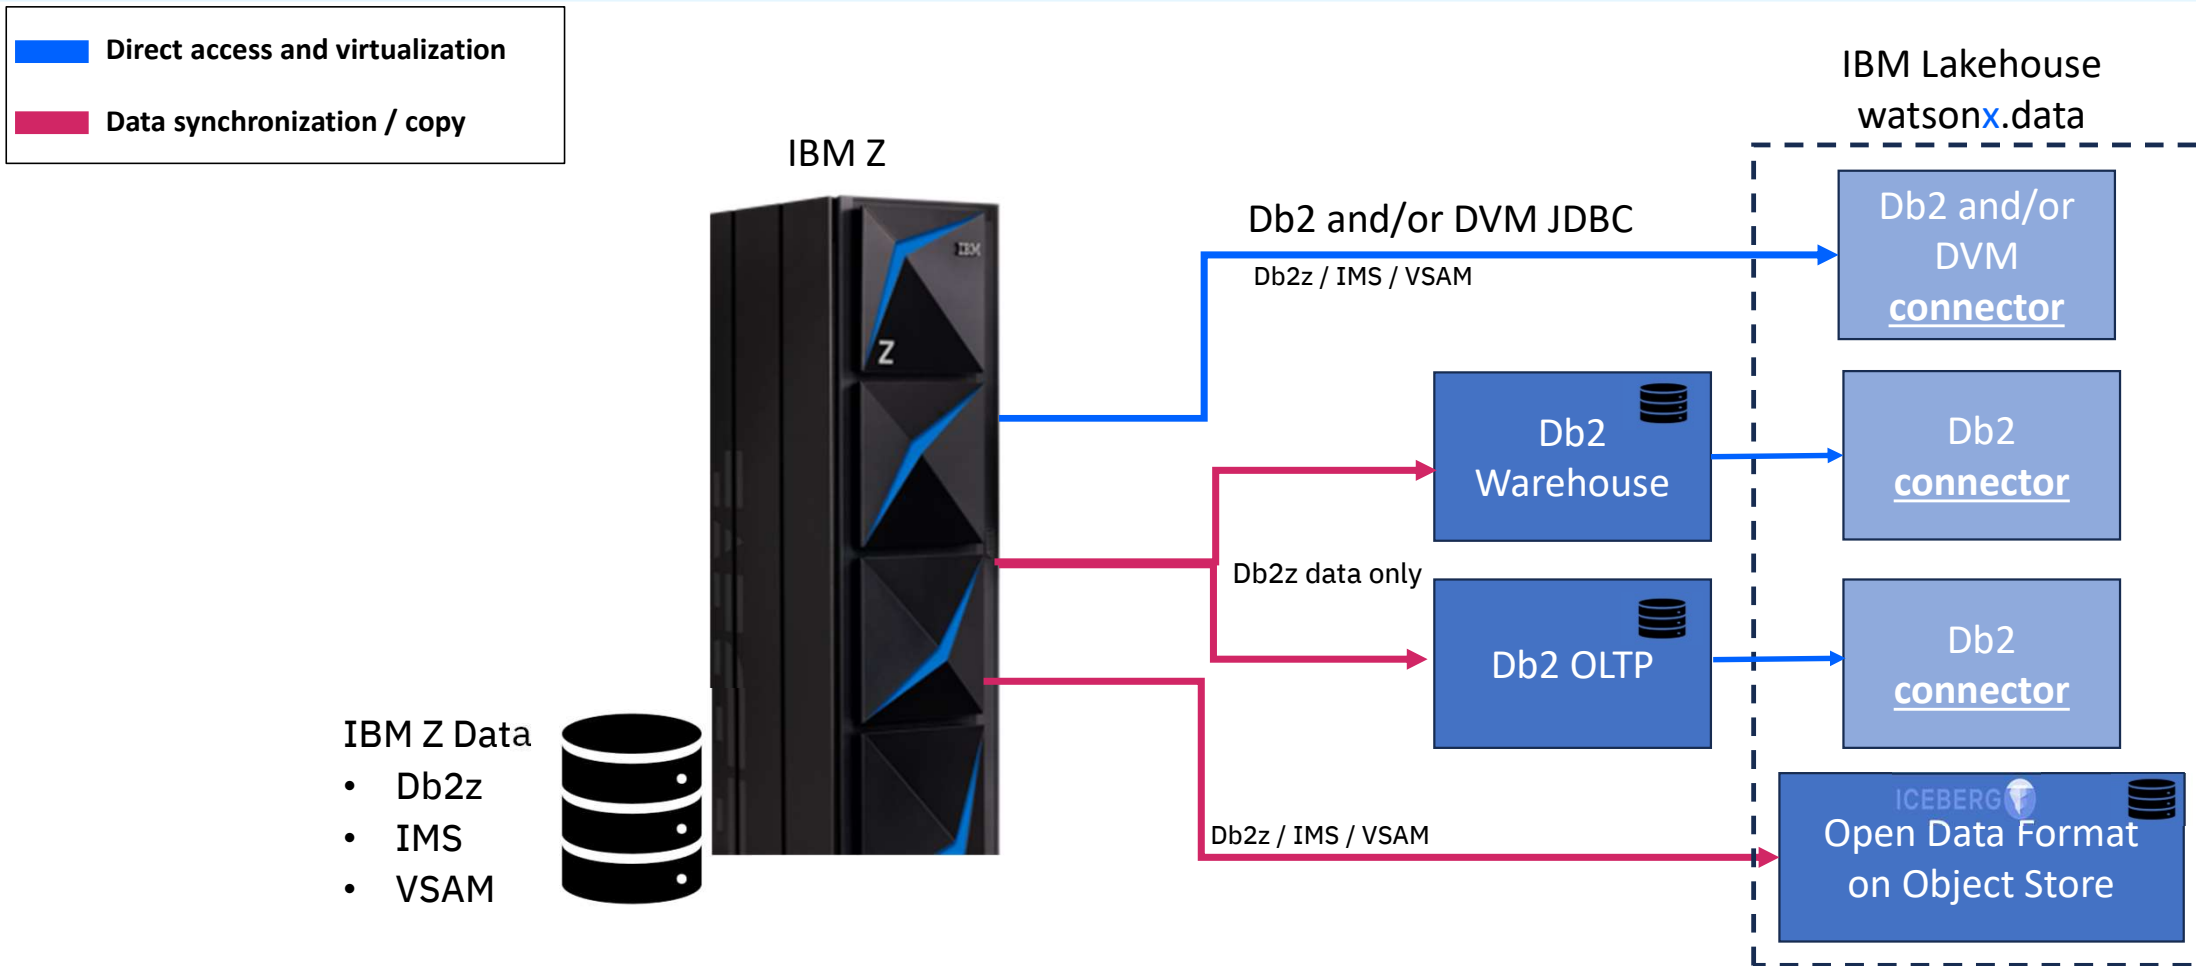
## Operational Data (read/write)



## Data Synchronization (unidirectional)



# Z Data Integration Options to Watsonx.data



DVM : Data Virtualization Manager for z/OS

- Virtualize (access) to live Db2, VSAM, IMS data in-place on IBM Z

# Major differences Data Gate for watsonx and CDC

	Data Gate for watsonx	IBM Change Data Capture (CDC)
<b>Sources</b>	Db2 for z/OS, IMS, VSAM only	Most major transaction processing environments including Db2 for z/OS/LUW/I, VSAM, IMS, Oracle, SQL Server, PostgreSQL incl. Amazon RDS, and Azure Database for PostgreSQL
<b>Targets</b>	Apache Iceberg in watsonx.data only	Many including databases such as Db2 for z/OS/LUW/i, VSAM, IMS, Oracle, SQL Server, PostgreSQL , warehouses such as Db2 WH, Google Big Query, datalakes, Kafka, files, ETL, public clouds
<b>Main capability</b>	Synchronize z/OS “table” as-is 1:1 to Apache Iceberg table in watsonx.data	Replicate selected tables (all or subsets of rows and/or columns) with filtering and transformation capabilities

watsonx.data & Data Gate for watsonx

# Collaboration with IBM research

- We found 2 inefficiencies in the OpenSource iceberg reference implementation (Apache Iceberg)
  - Writing data
  - Creating new snapshots (“commit”)
- Teamed up with IBM Research who co-created the ingest library with us that strongly optimizes the above codepaths.
- **NOTE: only testing ingest library locally!**

CPU: Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz  
Mem: 384GB  
S3: Minio on the same machine  
Hive Metastore: Running in non-authenticated mode on the same machine

Total rows inserted: 32M  
Rows per transaction: 1M  
Row size: 128 Bytes.

Rows Committed	Time for Spark-SQL(s) (single executor)	Time for Presto-SQL (s)
5k	26	8.85
100k	60	97.14
500k	84.35	129.5

# Collaboration with IBM research

- We found 2 inefficiencies in the OpenSource iceberg reference implementation (Apache Iceberg)
  - Writing data
  - Creating new snapshots (“commit”)
- Teamed up with IBM Research who co-created the ingest library with us that strongly optimizes the above codepaths.
- **NOTE: only testing ingest library locally!**

CPU: Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz  
Mem: 384GB  
S3: Minio on the same machine  
Hive Metastore: Running in non-authenticated mode on the same machine

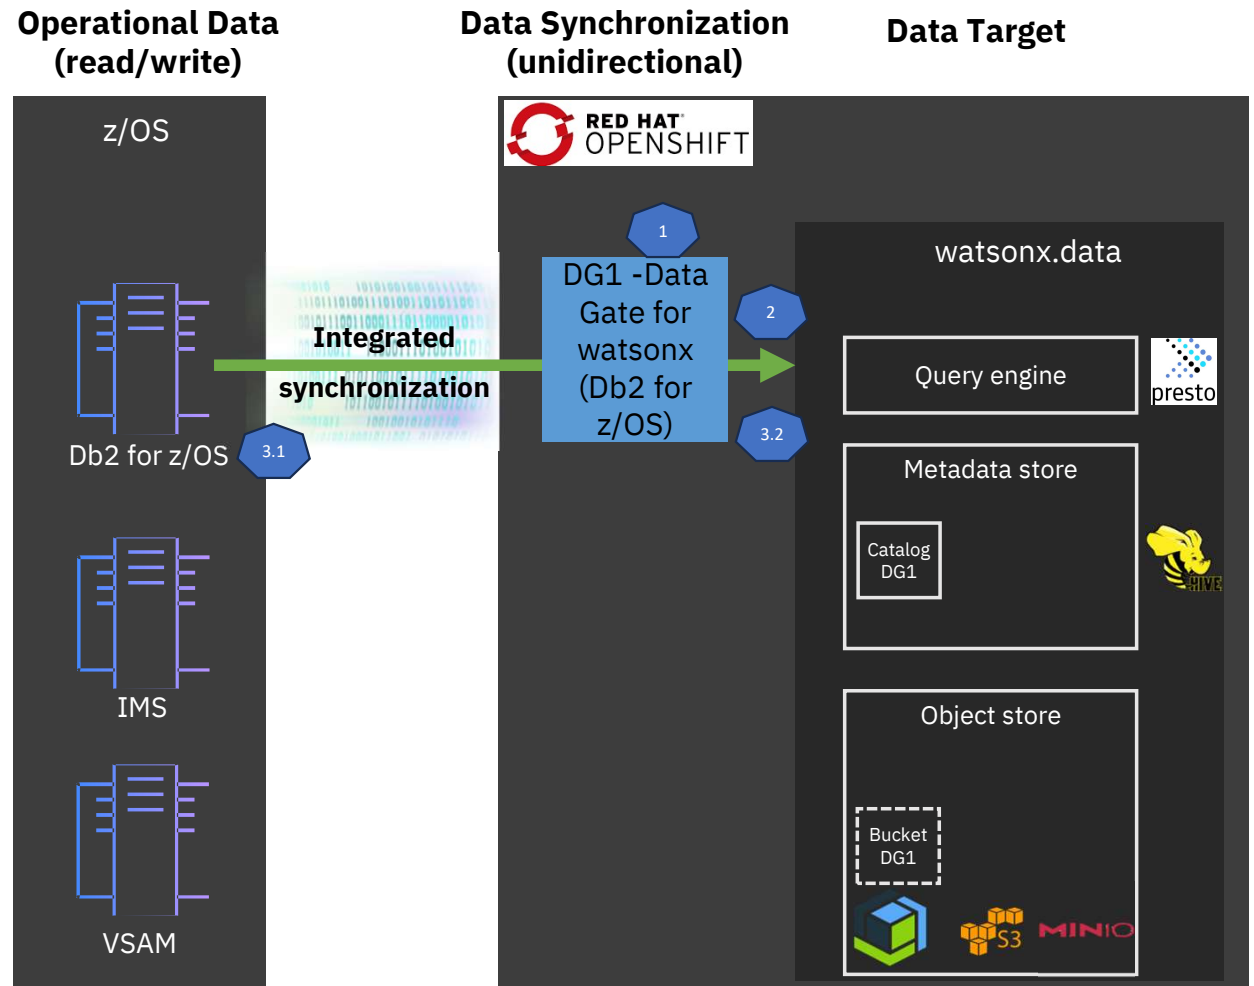
Total rows inserted: 32M  
Rows per transaction: 1M  
Row size: 128 Bytes.

Rows Committed	Time for Spark-SQL(s) (single executor)	Time for Presto-SQL (s)	Time for Tundra (s) (Single Instance)
5k	26	8.85	1
100k	60	97.14	1
500k	84.35	129.5	1

Up to **84x faster than Spark** and **130x faster than PrestoDB** when inserting 32M rows in 500k-sized transactions into an Iceberg file on object storage

# Data Gate for watsonx – Db2 for z/OS flow

- No installation necessary on Db2 for z/OS
    - only AT/TLS setup
    - Same mechanism and software used by IDAA
- 1) Add tables in Data Gate
  - 2) Load data through efficient Data Gate load feature
    - Unload Db2 for z/OS partitions in parallel
    - Load partitions in parallel in Iceberg
  - 3) Replicate data with Integrated synchronization
    - 3.1) Built-in log reader captures Db2 for z/OS changes
    - 3.2) Data Gate writes data into Iceberg tables in watsonx.data - High performance ingest



# Data Gate for watsonx – VSAM flow (IMS similar)

- Simple install and operation
  - Access Server and Apply components managed by Data Gate
  - User interacts with Data Gate UI or API
- AT-TLS setup required

## 1) Non-relational to relational mapping on the host side in Classic Data Architect

- 1.1) Define copybook
- 1.2) Create table mapping
- 1.3) Generate DDL model

## 2) Add tables through Data Gate UI / API

## 3) Refresh Bulk Apply (Load)

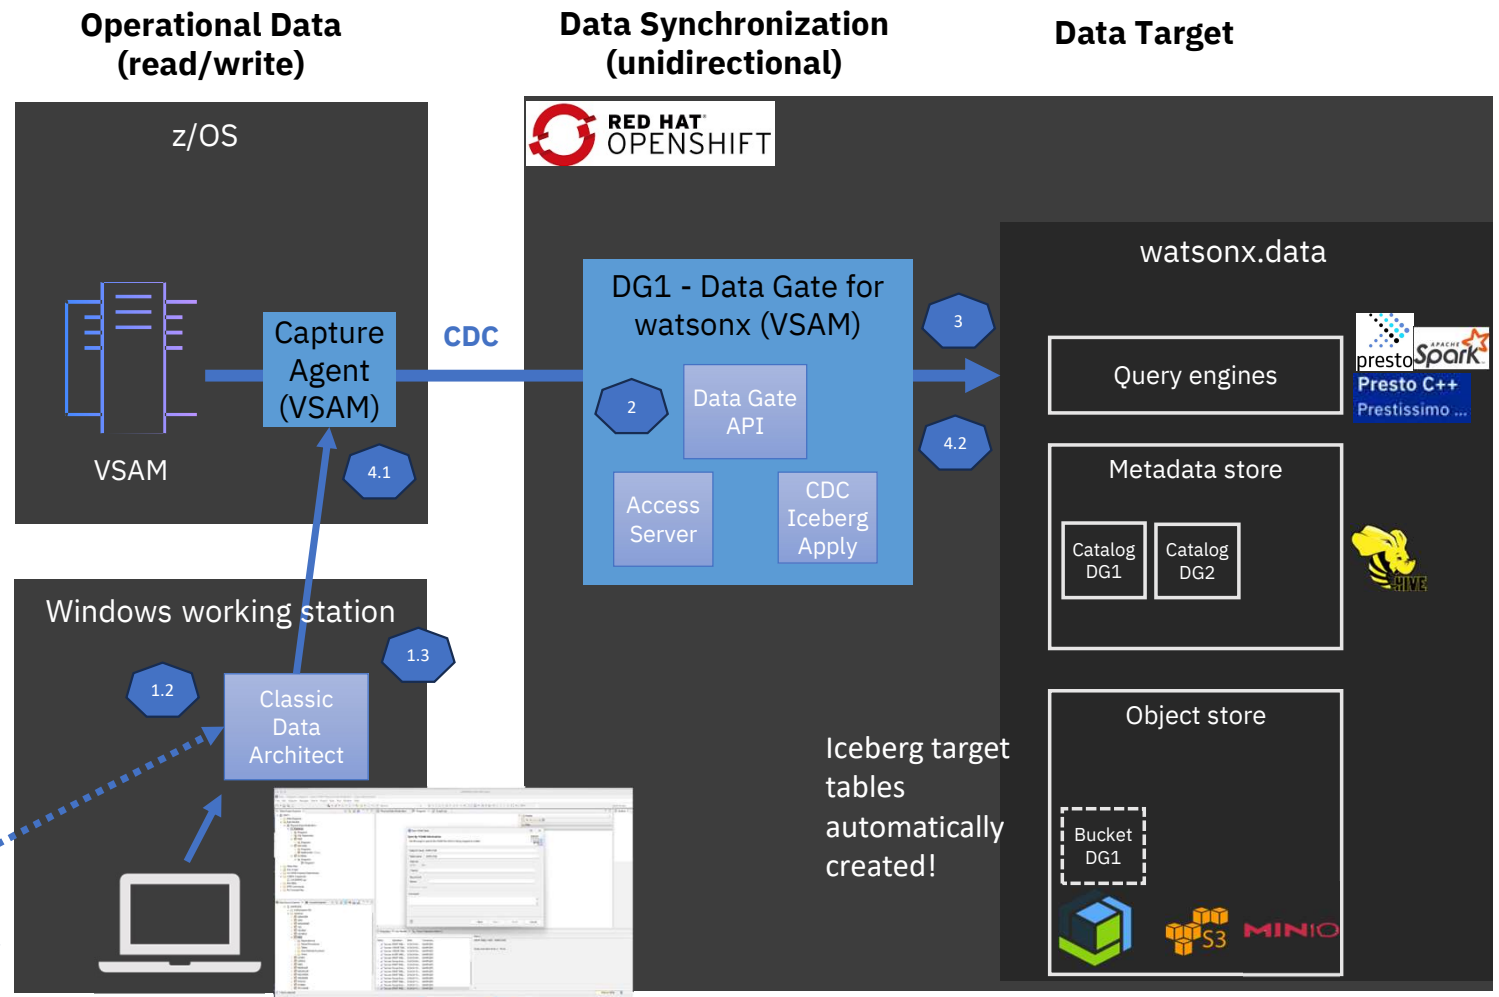
## 4) Mirror Bulk Apply (Replication)

4.1) CDC Capture Agent captures VSAM changes

4.2) Data Gate writes data into Iceberg tables on object storage (watsonx.data)

```

*****
* SAMPLE - EMPLOYEE COPYBOOK *
*****
01 EMPLOYEE.
05 ENAME PIC X(20).
05 PHONE PIC 9(9) COMP.
05 FILLER PIC X.
05 MAILID PIC X(6).
05 SALARY PIC S9(5)V99 COMP-3.
05 JOBID COMP-1.
05 EMPID PIC 9(9) COMP.
05 DEPTID PIC 9(4) COMP.
05 FILLER PIC X(2).
05 DEPARTMENT PIC X(15).
    
```



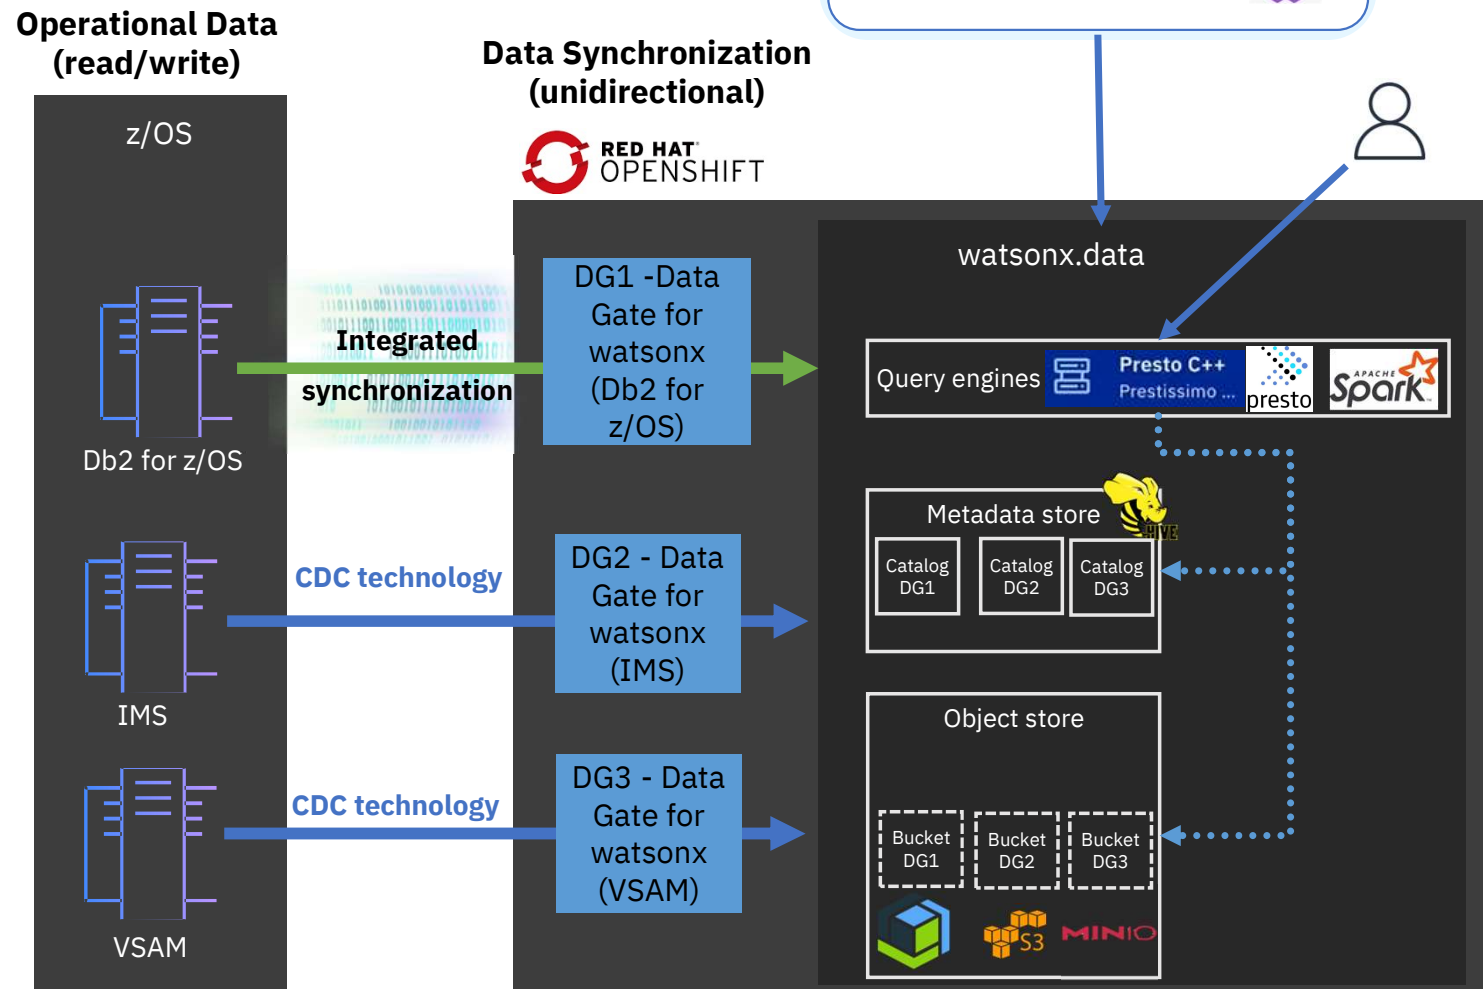
Iceberg target tables automatically created!



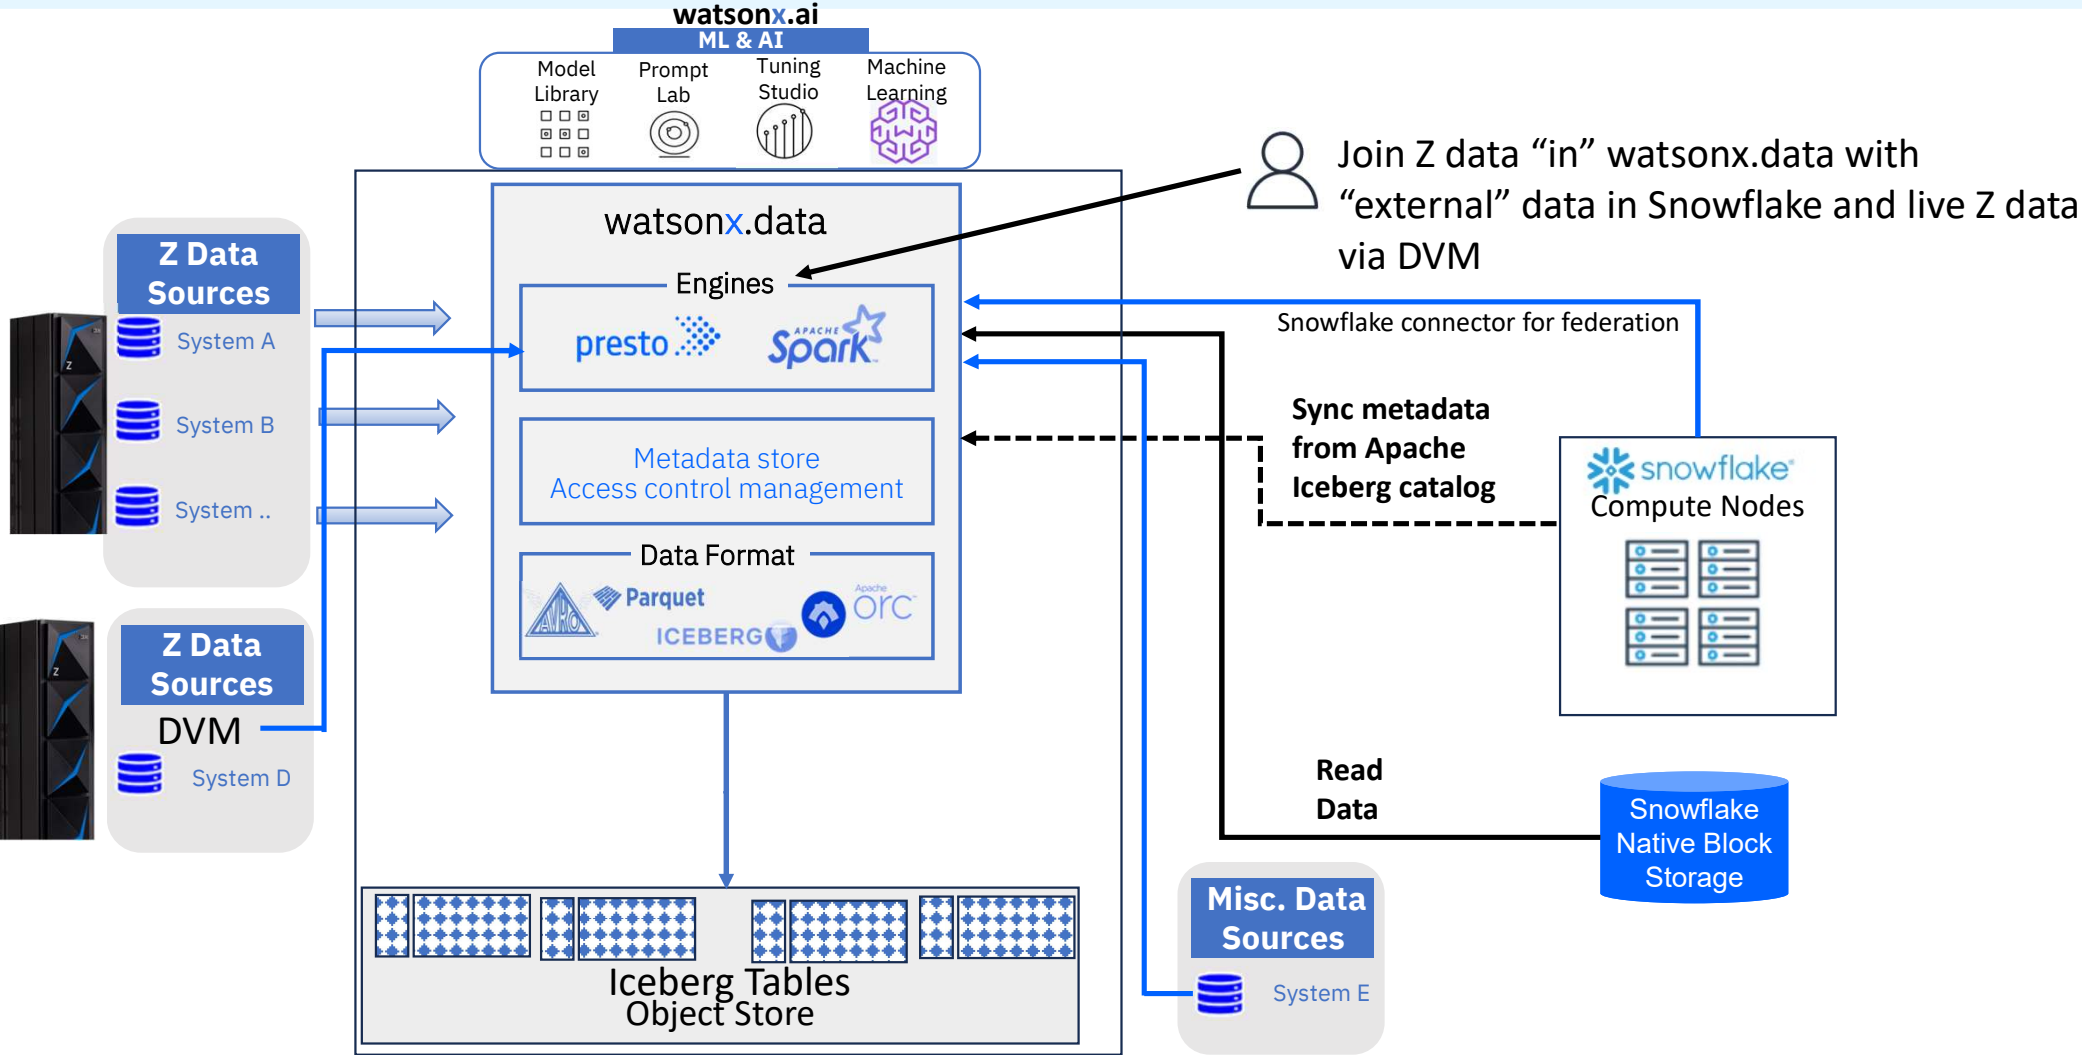
## Examples Data Gate for watsonx use cases

# Analytics and AI in watsonx.data

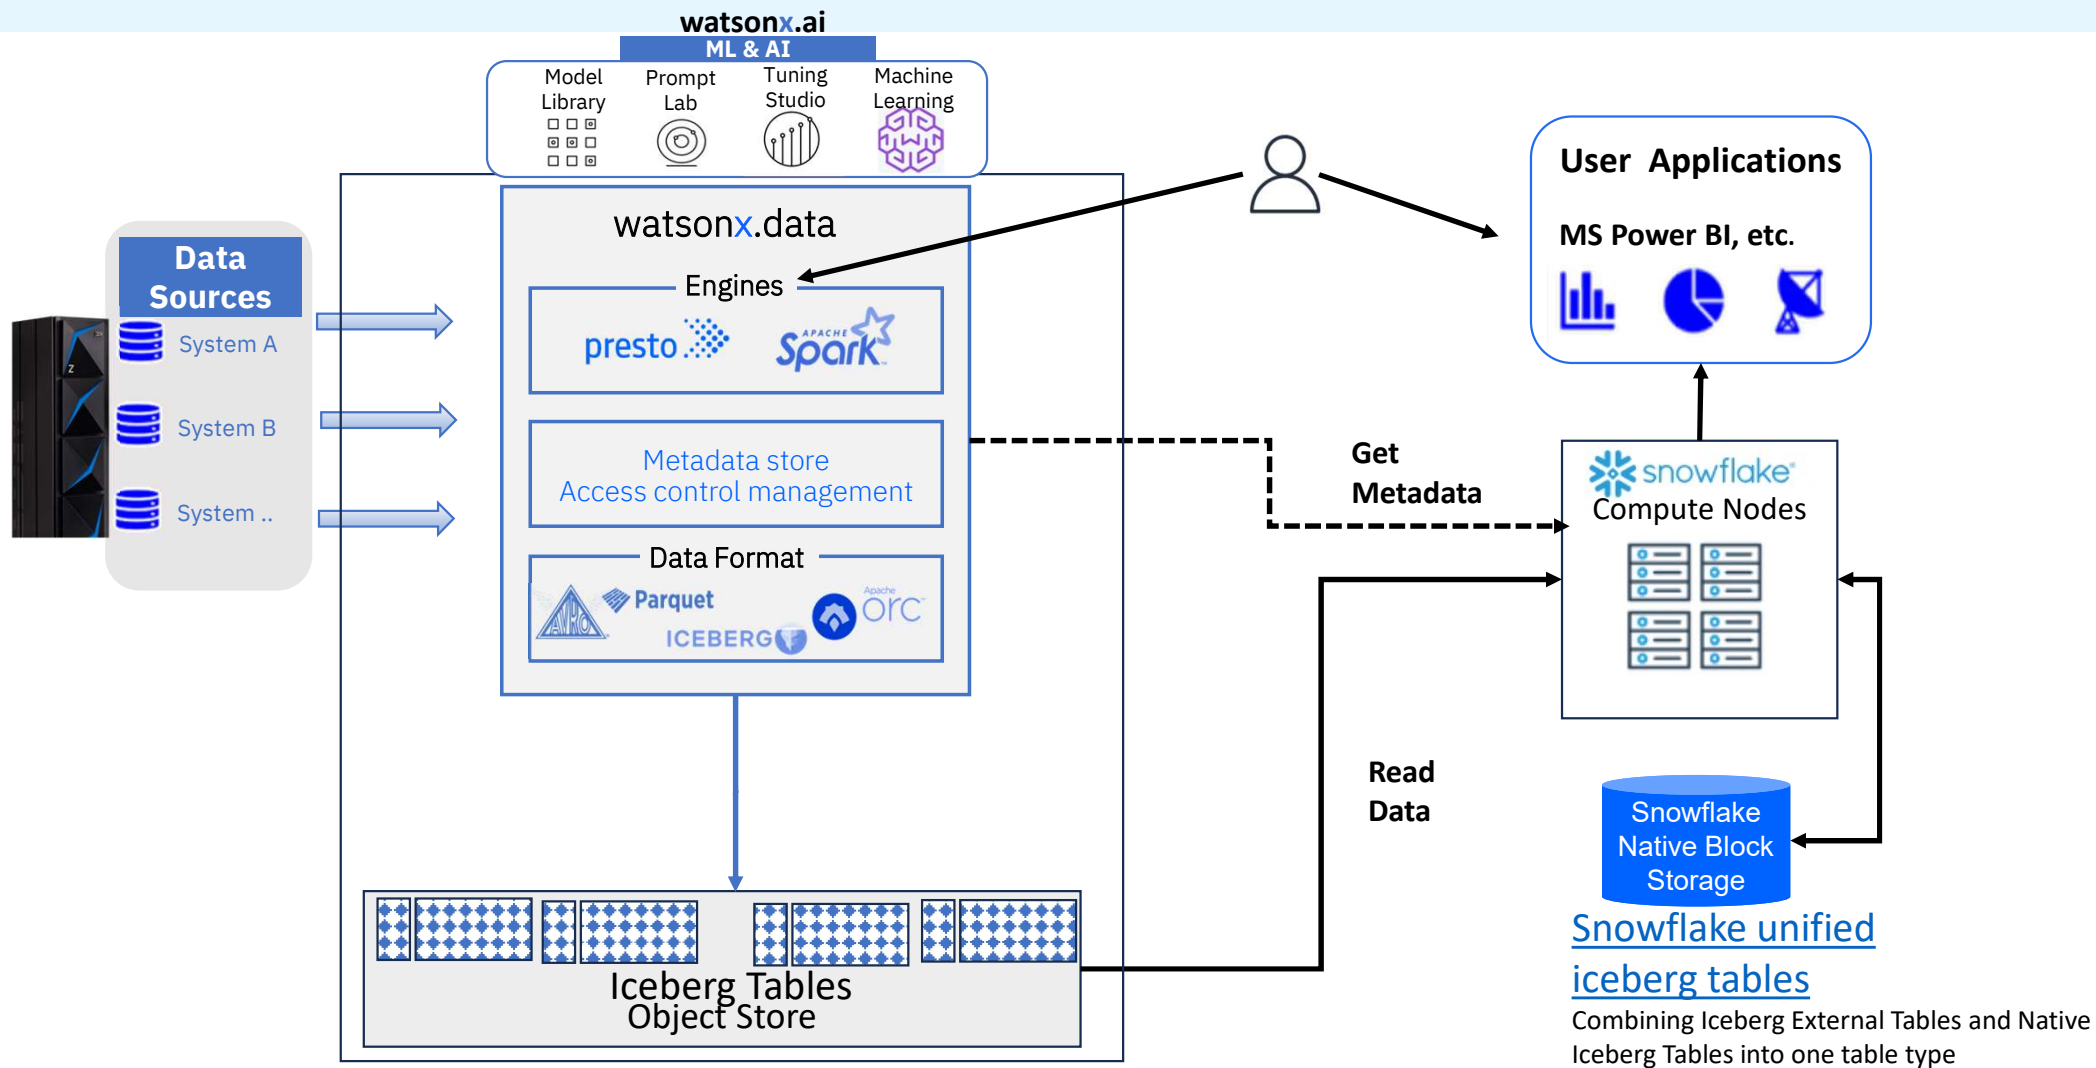
- Synchronize Db2 for z/OS, IMS, VSAM data to watsonx.data
- Run analytics on the same platform with Presto, Spark, etc.
- Highest replication performance because of co-location
- End-to-end encryption
- High Performance ingest into Iceberg
- Single end-to-end integrated platform
- Containerized, OpenShift-based installation (for non z/OS parts). Runs where watsonx.data runs.



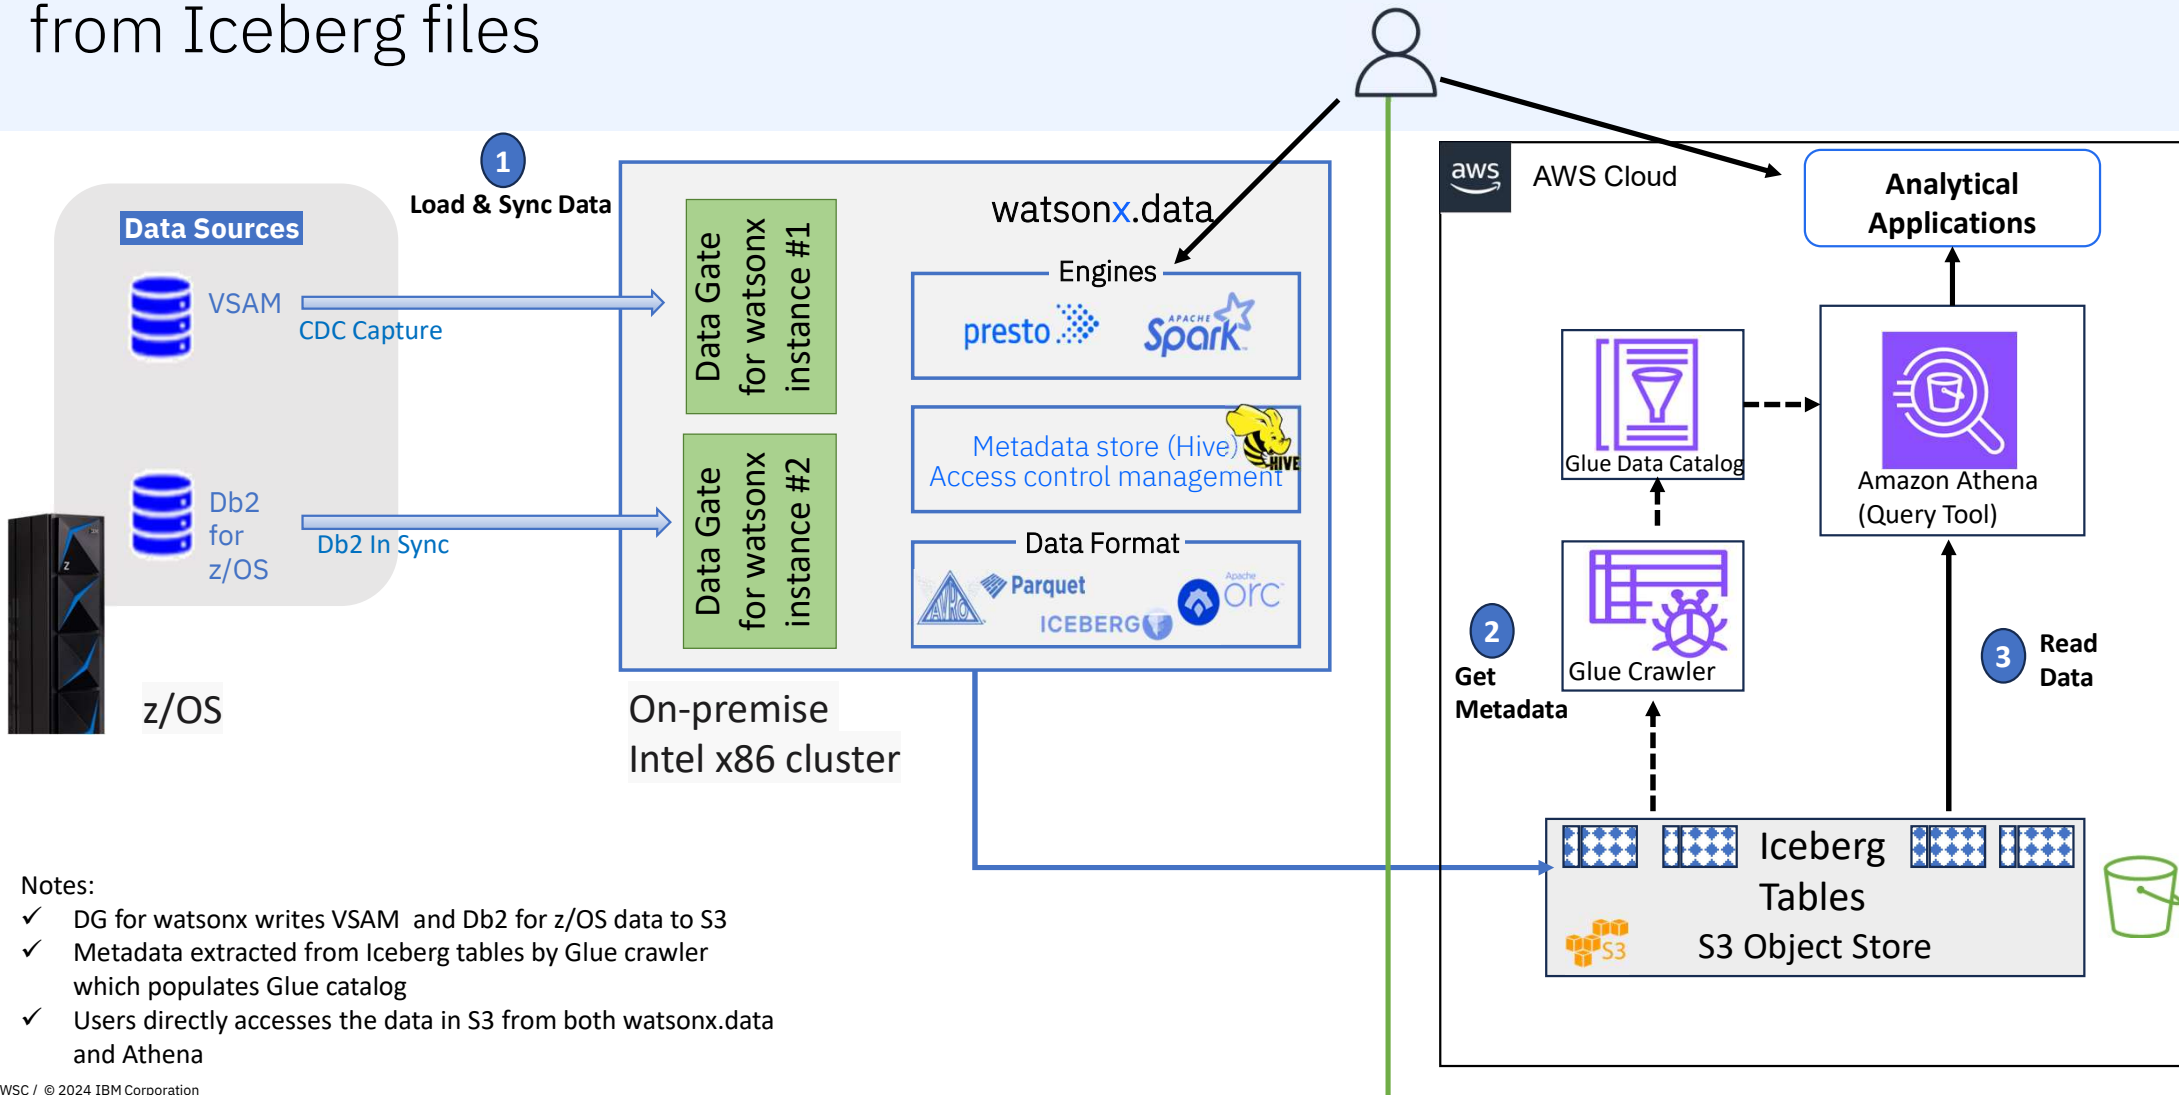
# Extending watsonx.data with existing Iceberg tables that are outside watsonx.data and live Z data



# Augmenting existing warehouse with Z data with using metadata in watsonx.data



# Customer PoT: Mainframe data in S3 in AWS and extracting metadata from Iceberg files



**Notes:**

- ✓ DG for watsonx writes VSAM and Db2 for z/OS data to S3
- ✓ Metadata extracted from Iceberg tables by Glue crawler which populates Glue catalog
- ✓ Users directly access the data in S3 from both watsonx.data and Athena

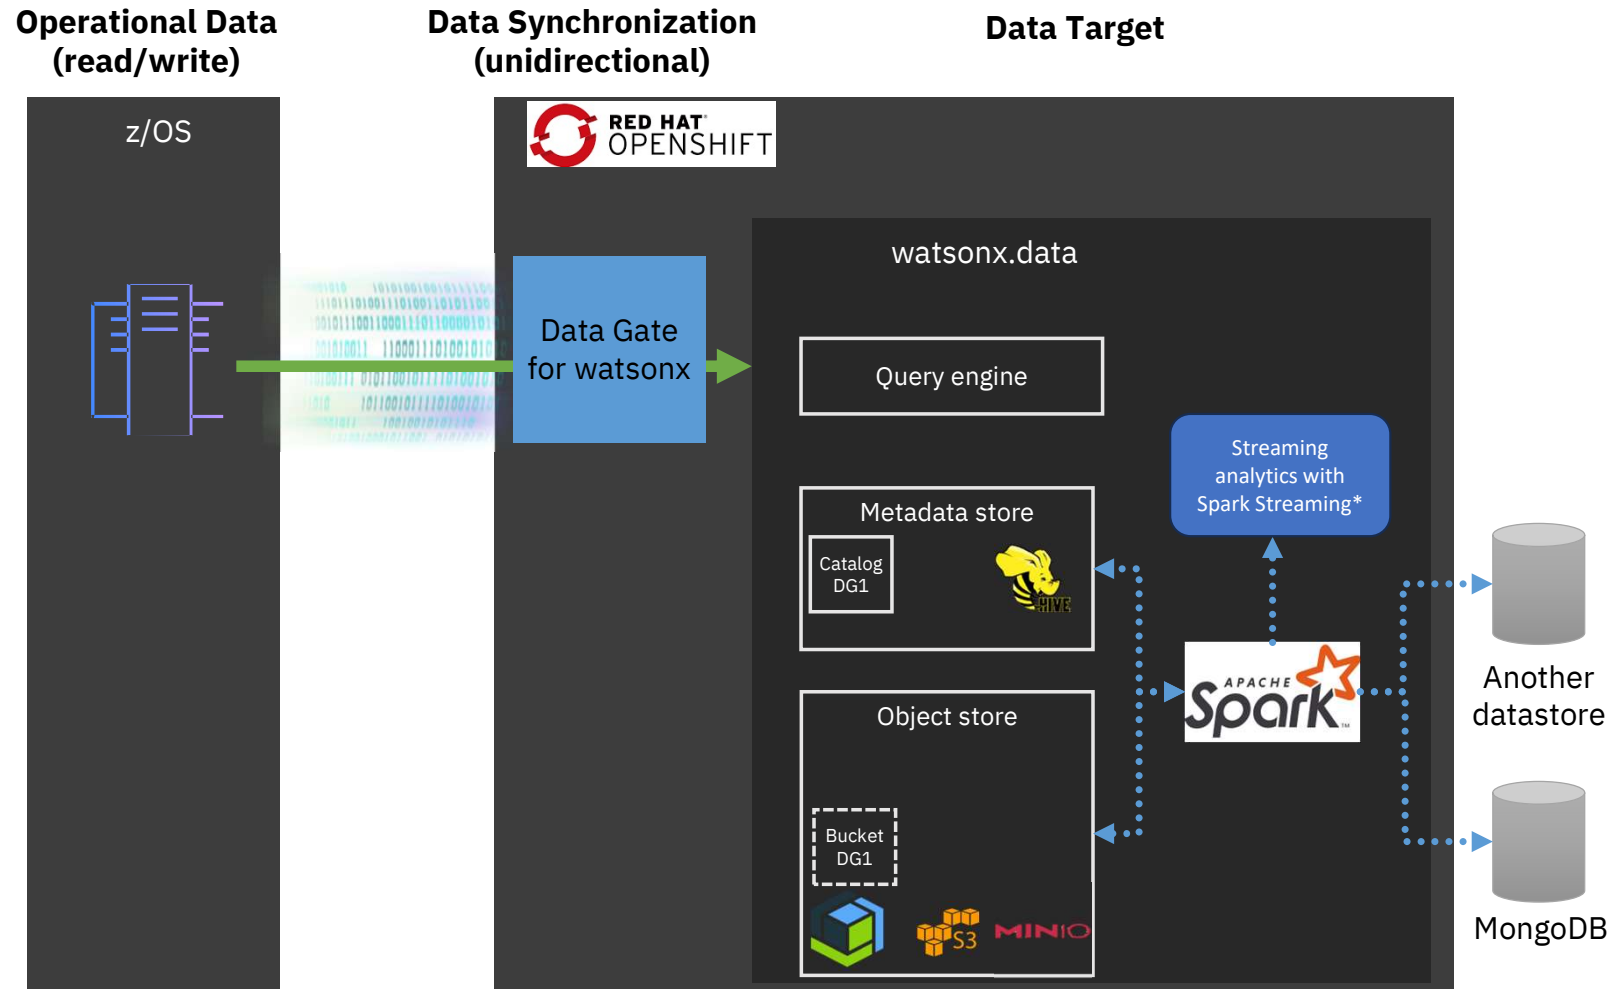
# Streaming with Spark Structured Streaming

- Synchronize Db2 for z/OS, IMS, VSAM data to watsonx.data
- Use Spark Streaming to
  - Run streaming analytics on the mainframe data
  - Write data to further data stores
    - As replacement for complex Kafka architectures

[Streaming Iceberg Table, an Alternative to Kafka?](#)



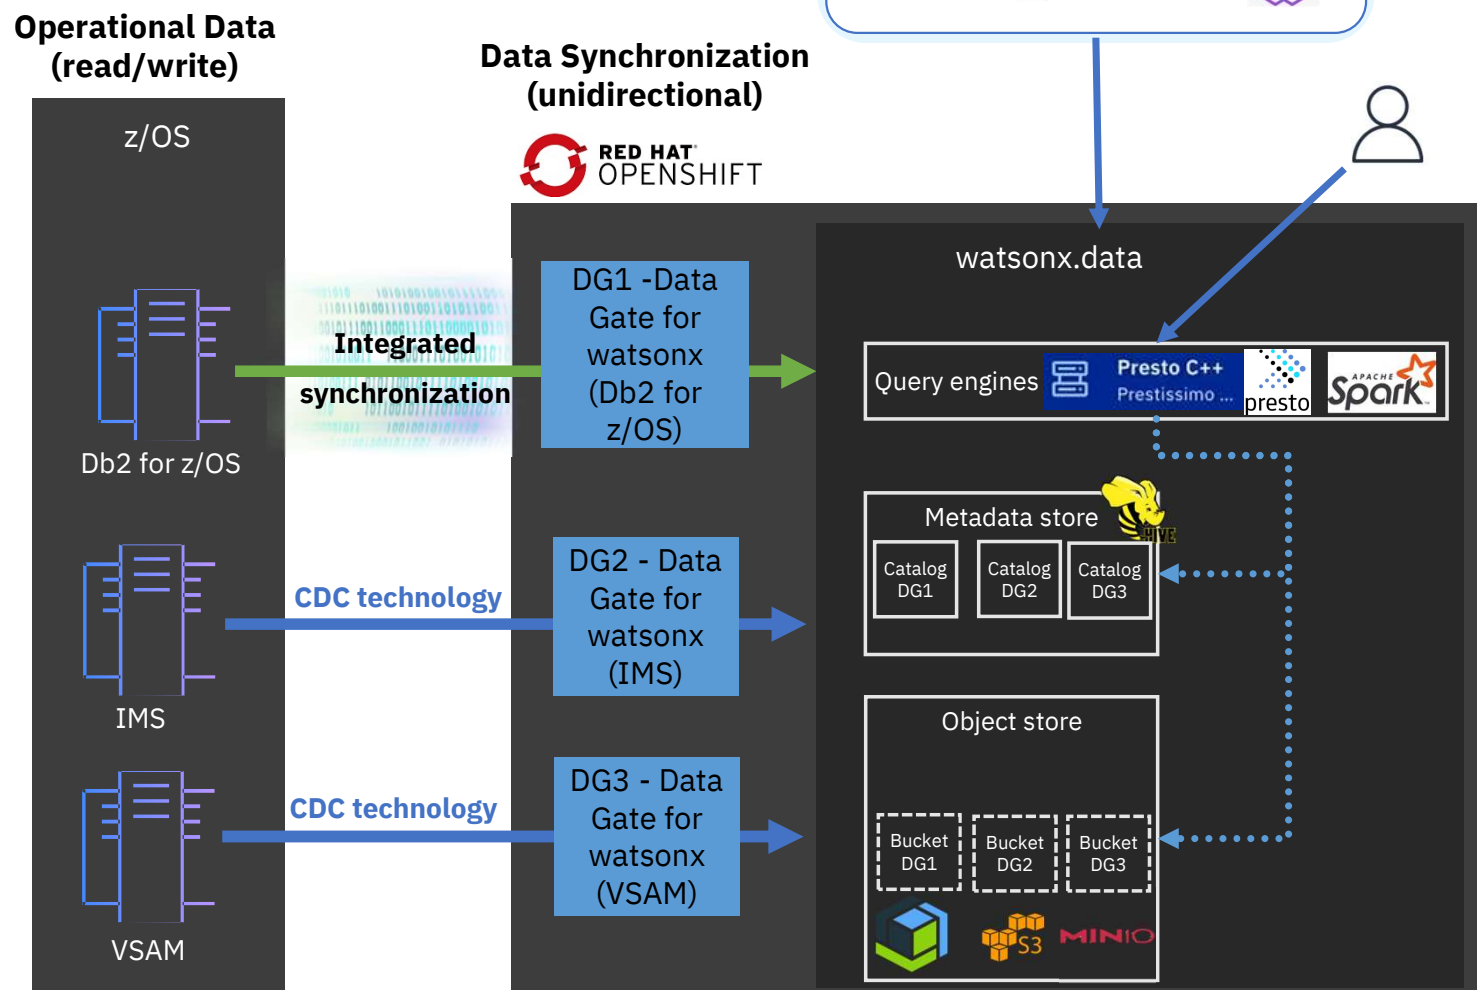
Integration with **StreamSets** planned for Q42024!  
"An end-to-end platform for smart data pipelines"



# Summary

# Leverage your mainframe data for AI and analytics by landing Db2 for z/OS, IMS and VSAM data directly in Iceberg tables on object storage

- Synchronize Db2 for z/OS, IMS, VSAM data to watsonx.data
- Run analytics on the same platform with Presto, Spark, etc.
- Highest replication performance because of co-location
- End-to-end encryption
- High Performance ingest into Iceberg
- Single end-to-end integrated platform
- Containerized, OpenShift-based installation (for non z/OS parts). Runs where watsonx.data runs.





Questions?