

# Db2 Analytics Accelerator and Db2 Data Gate

Present and Future

—

Daniel Martin  
Director and CTO for IDAA, Db2 Data Gate  
and OpenPages



IBM

# Legal Disclaimer

© IBM Corporation 2021. All Rights Reserved.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

# Daniel Martin

At IBM since 2010. Held various engineering and architecture roles contributing to Db2 LUW, Db2 z/OS, IDAA, Data Gate.

Spend most time as IDAA developer and architect. One of the original developers of the integrated synchronization technology.

<https://www.linkedin.com/in/daniel-martin-4a3a0998/>



## Daniel Martin

Director and CTO

Stuttgart Region · [Contact info](#)

# Agenda

Data Fabric with Z Data – 15'+

IBM Db2 Data Gate – 20'

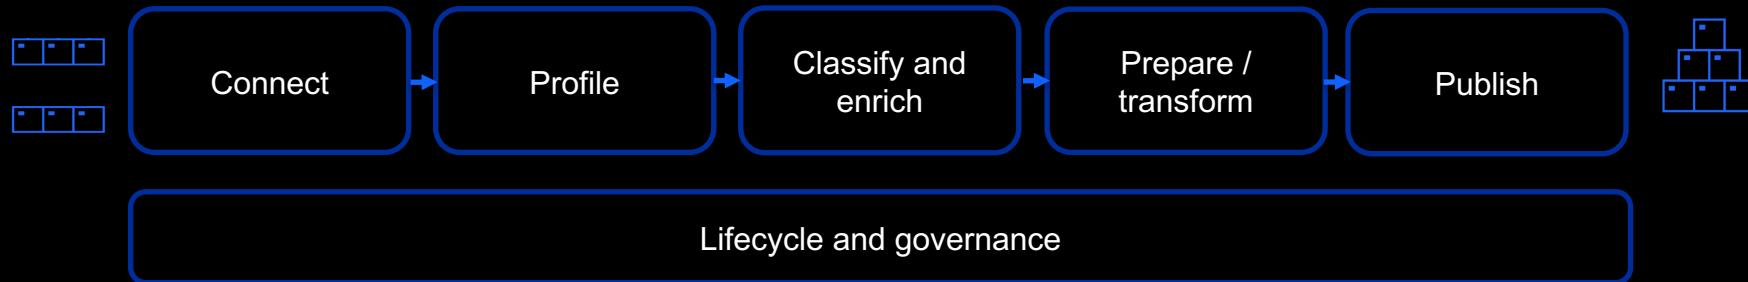
IBM Db2 Analytics Accelerator – 20'

# Data Fabric

# Data Fabric enables a “factory” approach for Data

Data Assets

Data Products



## Connect

Establish connectivity to physical data sources.

## Profile

Assess the quality of data assets.

## Classify and enrich

Classify data assets, assign data policies and rules, and enrich with semantics.

## Prepare / transform

Engineer data assets into trusted data products.

## Publish

Publish data products and enable subscription to data products based on **data contracts**.

## Lifecycle and governance

Implement DataOps principles throughout the lifecycle and enforce governance end-to-end.

# “Shop for Data” Experience for Data Consumers

**Namik** – *Reporting User*: Wants to build an interactive dashboard to look at our Card users

**Cuneyt** – *Analytics User*: wants to use his Python skills to do linear regression and learn about Card usage trends




**Eberhard** – *Data Scientist*: wants machine learning to improve Card fraud detection

*How do they learn about data assets that exist?*

*How can they use data products I created for them or reuse existing data products that may already fit their needs?*



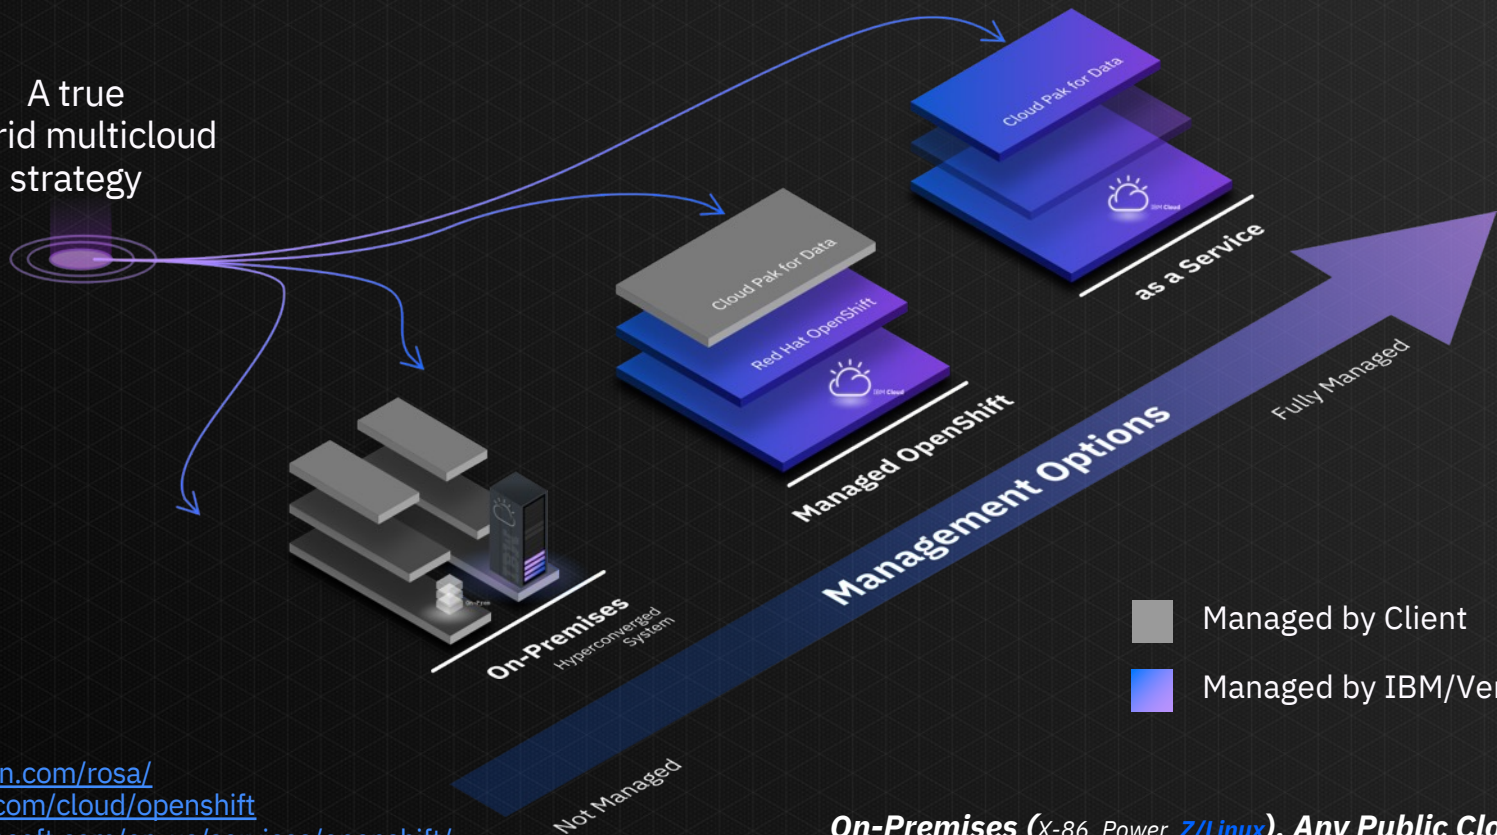
Data consumers



Data engineer for Credit Card Domain

# Deploy Anywhere.

A true hybrid multicloud strategy



Managed OpenShift

ROSA: <https://aws.amazon.com/rosa/>

ROKS: <https://www.ibm.com/cloud/openshift>

AROS: <https://azure.microsoft.com/en-us/services/openshift/>

**On-Premises (X-86, Power, Z/Linux), Any Public Cloud**

# The Metadata Catalog Plays a Central Role

## Enables

**Discovery and search** through metadata, tags, terms and categories

**Governance** through audit logs, lineage, rule enforcement (e.g. data masking depending on user or location)

Supports **decision** making for most **suitable data integration method** (see later)

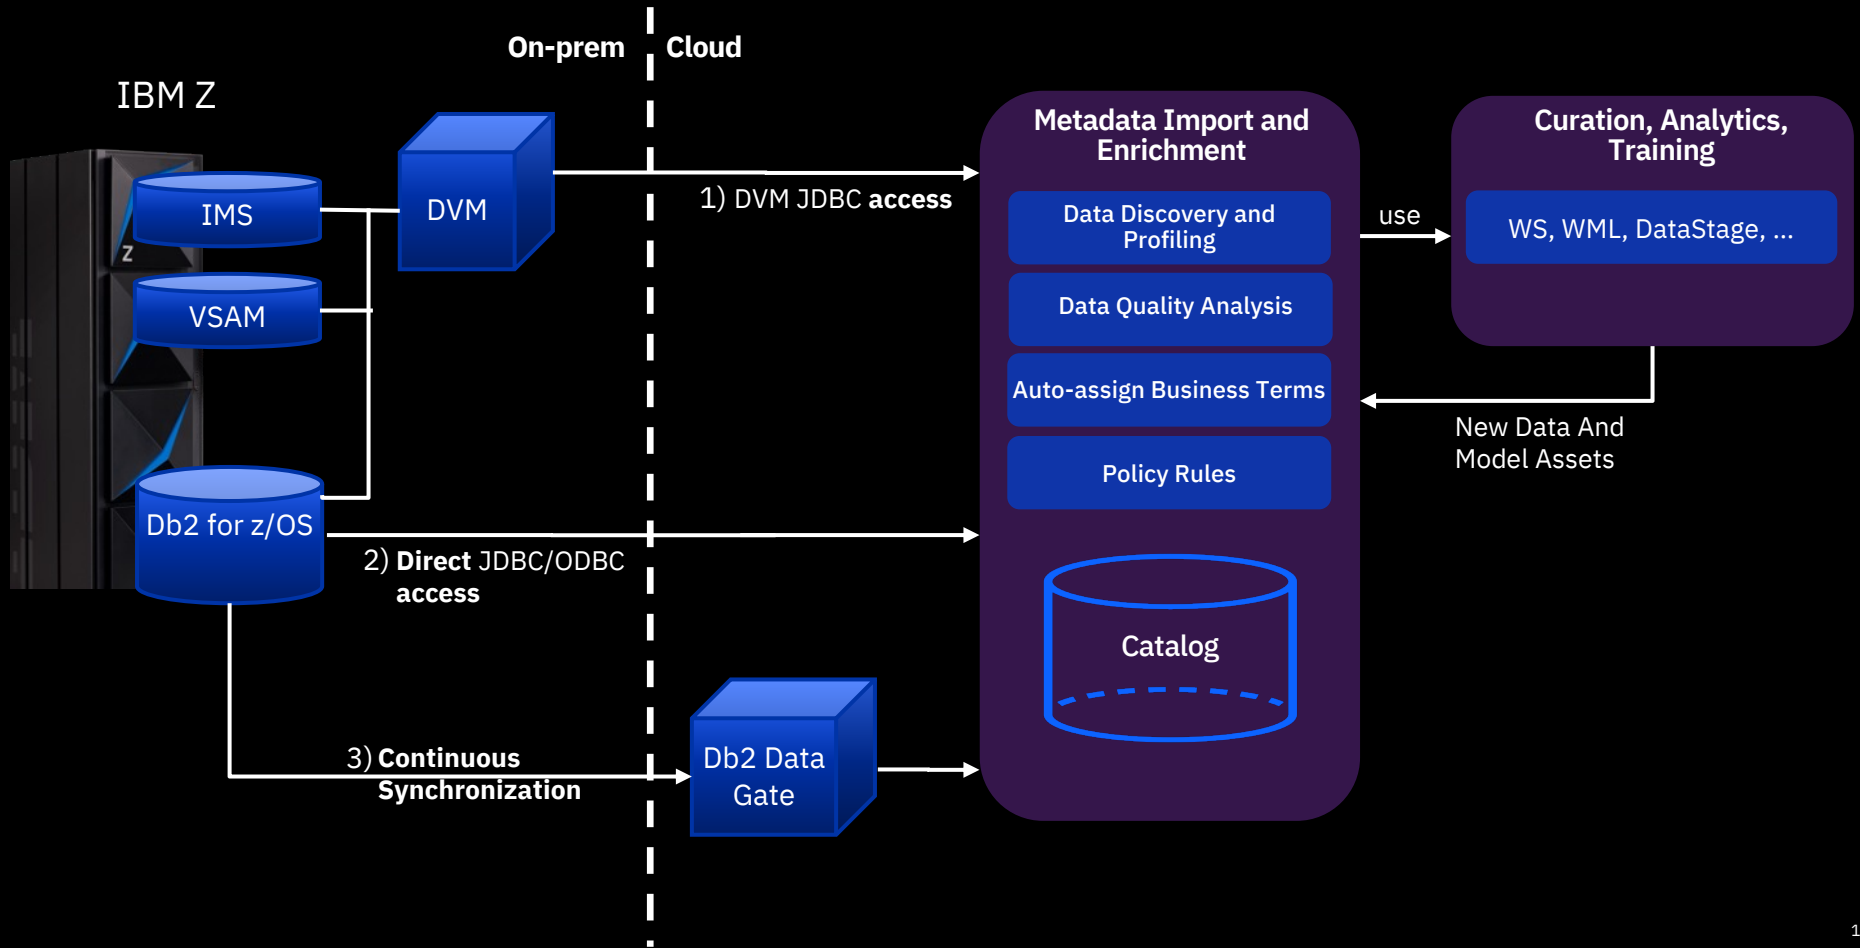
**Latest CP4D v4.5** release enables all these use-cases for mainframe data sources:

Db2 z/OS with **direct connection** (JDBC)

Other, **non-relational sources** through **Data Virtualization Manager** (DVM)

**IBM Db2 DataGate**  
Starting with 4.5, **DataGate can synchronize data** (to Db2) and **metadata** (to catalog)

# Integration Options for IBM Z Data with Data Fabric



# Data Gate

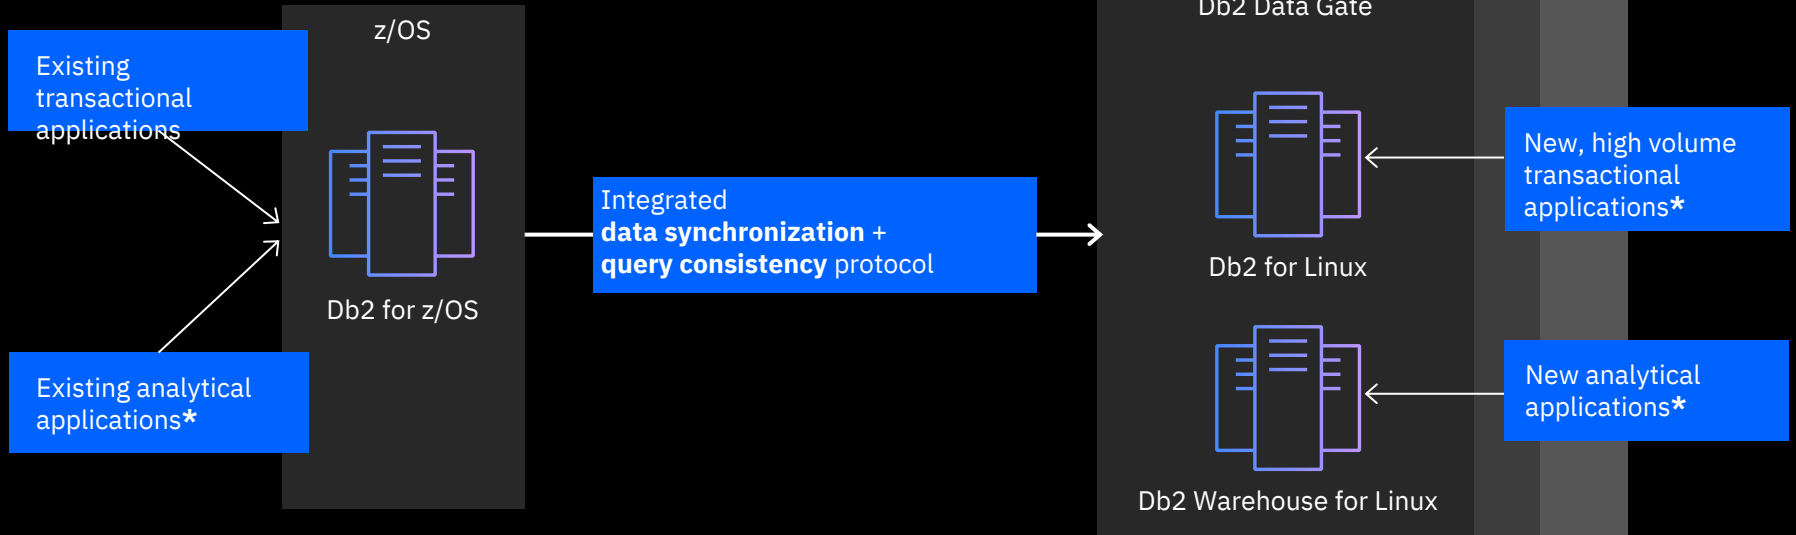
# Db2 for z/OS Data Gate

## Key values

- **Simplicity** - Data Gate offers a radically *simplified user experience for the task of creating replicas* of Db2 for z/OS tables and keeping them transactionally consistent.
- **Isolation** – Cloud Pak for Data tools and applications can access zSystems data *without any impact to the transactional workload* running on zSystems. **Enabling self-service for consumers.**
- **Low latency** - An integrated, patented synchronization protocol maintains data currency.
- **Cost-effective** - Db2 data changes capture processing is *fully zIIP-enabled* to ensure the lowest processing cost on zSystems.



# Data Gate Technical Overview



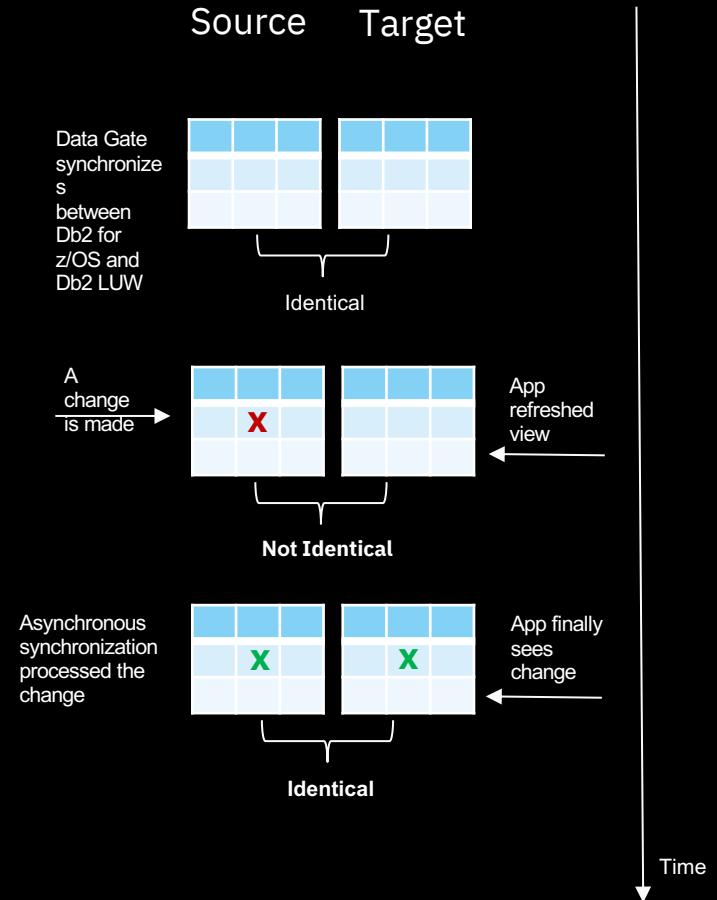
\* Optional SQL syntax to *guarantee* returning most recently committed data

1. New SQL syntax on Db2 LUW for queries referencing Db2 for z/OS tables: **guarantee query to return most-recently committed changes.**
2. **Query acceleration** support (IDAA use-case): extend Db2 for z/OS with Cloud resources to transparently **accelerate** analytic queries **or archive** Db2 z/OS tables.
3. **Data Fabric** integration: **synchronize data and metadata** of Db2 z/OS tables to enterprise catalog to enable “shop for data” experiences, data governance and lineage.

# 1. New SQL Syntax for Transactionally Consistent Query Results

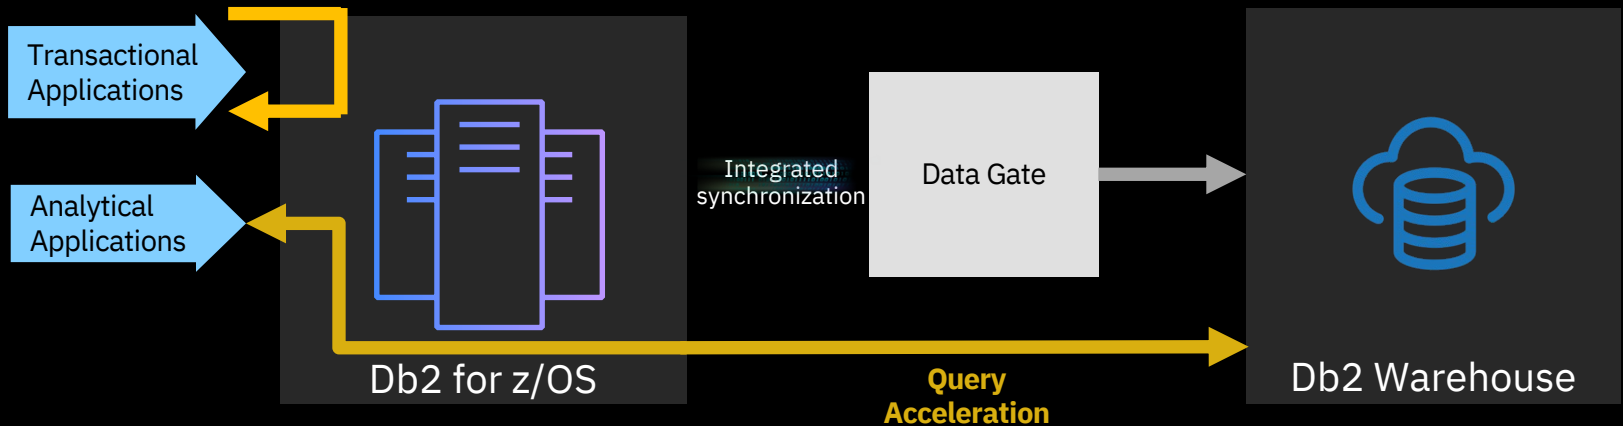
## Scenario

1. Data Gate is setup to synchronize tables between Db2 for z/OS (source) and Db2 LUW (target)
2. Without resultset consistency guarantee (new, optional SQL syntax in Db2 LUW) there is no control over what data is seen on the target. **Recent changes may or may not be included in queries**
3. With the new SQL syntax which can be specified on a per-query level or on the subsystem level, Cloud applications get the guarantee to always operate on most-recently committed data wrt. Db2 for z/OS.

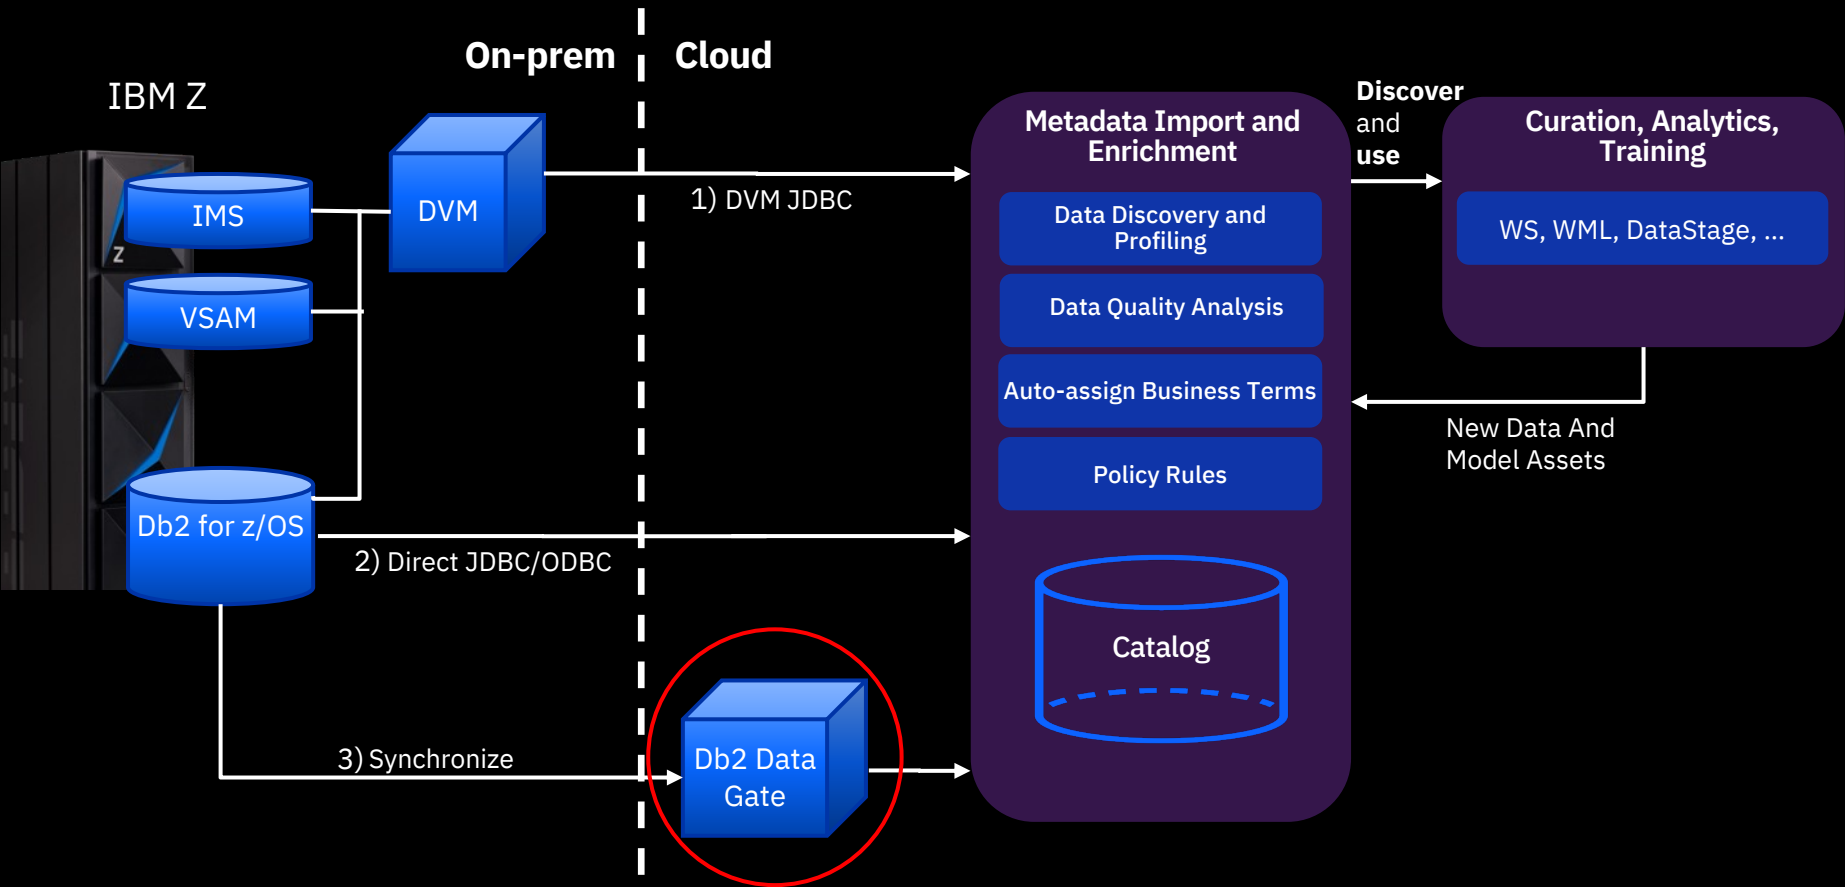


## 2. Db2 z/OS Analytical **Query Acceleration** and Archiving

- Accelerate data-intensive queries transparent to applications: Db2 z/OS optimizer decides if query should be routed to “accelerator” or stay in mainline Db2 (e.g. because of index access).
- Archive “cold” partitions of Db2 for z/OS tables to Db2 LUW but still keep them query-able.



# 3. IBM Z Data Integration with a Data Fabric



# Demo

(DataGate UI + Catalog Integration)

# Data Gate in Cloud Pak services catalog

Services catalog /

## Db2 Data Gate

IBM Enabled ✓

- Documentation
- Instances
- New instance

The screenshot displays the IBM Cloud Pak for Data console for the 'Data-gate-1' instance. It includes a navigation bar, a time filter set to '6 hours', and three main monitoring panels: 'Data-gate-1' status, 'Synchronization latency', and 'Avg synchronization throughput'. Below these is a 'Tables' section with a search bar and a table listing active tables.

Table	Schema	Table size	Sync	Last synced	Status
<input type="checkbox"/> breakfast table	528	1 MB	Disabled	3/23/20 5:29 AM	<input type="checkbox"/> Inactive

Extract, load, and synchronize mission-critical Db2 for z/OS data for high volume transactional or analytic applications.

### Version

2.0.2

### Description

The IBM Db2 for z/OS Data Gate service extracts, loads, and synchronizes your Db2 for z/OS data to IBM Cloud Pak for Data for quick access by your new, high volume, read-only transactional applications and analytic or AI workloads.

The service propagates your Db2 for z/OS data to a Db2 Warehouse or Db2 database on IBM Cloud Pak for Data. Through its high throughput and low latency synchronization technology, the service provides you near real-time access to the data without degrading the performance of your core business transaction engine.

Compared to custom-built solutions, Db2 Data Gate reduces the cost and complexity of application development and ensures that your data is current, consistent, and secure.

### Prerequisite services

- Feature dependencies
- Db2
- Db2 Warehouse

### About the developer

IBM




# Data Gate – Instance creation

Services catalog / Db2 Data Gate / Create data gate instance

Select target database type and instance. Create one if none exists on the cluster. Your Cloud Pak for Data userid is used to access the target database, which requires the admin role.

- Transactional
- Analytic [Create analytic database](#)

### Compute resources



### Namespace

zen

### Storage

Create new storage if no existing storage is available.

- Create new storage
- Use existing storage

managed-nfs-storage Size (GB) 50

### Route

Host dg2.apps.dgdemo7.cp.fyre.ibm.com Port 44443



# Data Gate – Pairing with Db2 for z/OS



- Define source    Select tables    Finish

## Source database

Host\*  Secure DDF port\*

Location\*

Username\*

Password\*

## Log reader

- Use source database connection information for the log reader
- Use different connection information for the log reader

## TLS certificate

Upload the certificate for connecting Db2 for z/OS and Db2 Data Gate servers\*

[Drop your PKCS12 file here or browse for a file to upload \(Maximum file size: 500 KB\)](#)

dgshare2021.p12

Keystore password\*



# Data Gate – Add Tables

Instances / Data-Gate-1

## Add table

Search and select tables for synchronization

Schema

Schema	Table / LOANDEMO
<input checked="" type="checkbox"/> LOANDEMO 5 tables	<input checked="" type="checkbox"/> ACCOUNT N/A
<input type="checkbox"/> LOANDEMO13 4 tables	<input type="checkbox"/> APP02 N/A
<input type="checkbox"/> LOANDEMO15 4 tables	<input type="checkbox"/> BORROWER_HISTORY N/A
<input type="checkbox"/> LOANDEMO2019 4 tables	<input checked="" type="checkbox"/> CREDIT N/A
<input type="checkbox"/> LOANDM2 4 tables	<input checked="" type="checkbox"/> ORDERS N/A

Tables selected Clear all

LOANDEMO	3 tables
<input checked="" type="checkbox"/> ACCOUNT	N/A
<input checked="" type="checkbox"/> CREDIT	N/A
<input checked="" type="checkbox"/> ORDERS	N/A

# Data Gate – Add Tables confirmation

## Add table

Confirm your selections or go back to make changes if needed.

Total schemas 1

Total tables 3

Total estimated table size N/A

Enable synchronization and load tables

# Data Gate Dashboard

## DC1A11

- Overview
- Monitor
- Table
- Query

### DC1A11

Target: Db2\_wh\_svt

Source: STEC2@idaadb2z.mycompany.com:691

#### Status

- Data gate server active
- Synchronization enabled  On
- Query routing disabled  Off

New!

### Data Catalog Integration

Status: ✔ Available [→](#)

Last Asset Synchronization: N/A

[Synchronize now](#)

### Table




Table: 26 total

Active	xxx	Pending	xxx
Loaded	xxx	Error	xxx
Loading	xxx		

### Query



Query: 35 total


Successful	xxx	Queued	xxx
Executing	xxx	Failed	xxx

# Data Gate Monitor View

- 1 hour
- 6 hours
- 12 hours
- 24 hours
- 3 days
- Custom

Time point:  
08/16/2022 17:33:43


### Synchronization latency



Synchronization latency

- Latency: 0.295

### Synchronization throughput




Synchronization throughput

- Throughput: 3700

### Source data and logs

- Inserted rows
- Updated rows
- Deleted rows
- Compensated rows
- Utility logs



# Product Roadmap

3Q 2022	4Q 2022	2023
<p><b>Data Gate 2.1</b></p> <p>Data fabric: Publish meta-data to Watson Knowledge Catalog with the press of one button</p> <p>Acceleration of Db2 for z/OS queries</p>	<p><b>Cloud Pak for Data 4.6</b></p> <p>WKC synchronization: automatic and asynchronous</p> <p>More monitoring options in UI</p> <p>Cloud Pak for Data 4.6 platform mandates</p> <p>CSV-Export of Data Gate instrumentation data</p>	<p><b>In planning. Candidates:</b></p> <p>REST-API to access replication state on a table level</p> <p>stunnel in a container image to simplify network connection</p> <p>Usage of GDPS as a disaster recovery solution on LinuxONE</p> <p>Watson Knowledge Catalog: provide Lineage</p>

# **Db2 Analytics Accelerator**

# Status of our Appliance Hardware Platform and the use of Infosphere CDC

- The Integrated Analytics System (IIAS, M4002 generation) inventory of hardware is exhausted. Therefore, the **End of Marketing** (not the End of Support!) is in effect since June 30, 2021.
  - End of Support for M4002 has not been announced but is likely in **September 2026** (5 years after EoM).
- We announced the **deprecation** of incremental update based on **IBM InfoSphere Data Replication (CDC)** with the availability of 7.5.5 and encourage our customers to phase-out the use of CDC for incremental update on the Accelerator and move to **IBM Integrated Synchronization**.

# Hardware Platform Roadmap for Db2 Analytics Accelerator

- Db2 Analytics Accelerator on **IBM Z** (deployed on IFL processors on IBM Z or LinuxONE) will be the **only deployment option** after the EoM of IIAS.
- Db2 Analytics Accelerator on IBM Z (running on IFLs or LinuxONE) is **not affected by the appliance EOM** in any way.
- Both platforms have a very active roadmap and continue to be enhanced and supported
  - Db2 Analytics Accelerator on IIAS will **benefit from the same enhancements** that flow into the deployment option on IBM Z
  - Except those that only relate to Z-platform specifics, e.g., GDPS or HMC



# Db2 Analytics Accelerator on IBM zSystems

## A fully integrated software appliance

- Running on IBM zSystem servers in a **Secure Service Container (SSC) LPAR**
- “download & go” – very easy to deploy
- Delivers the full SW stack including operating system, docker engine, Db2 Warehouse engine, Accelerator components, and infrastructure management

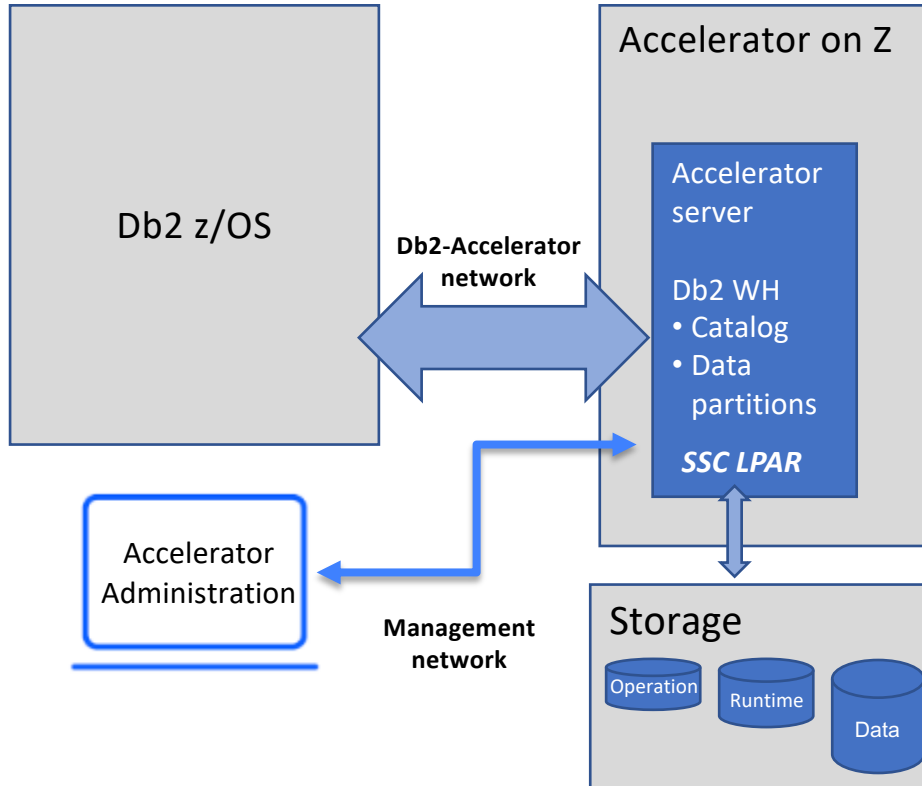
## A solution leveraging IBM zSystems infrastructure

- Runs on all IBM Z systems starting with z13 – **z16** or on LinuxONE
- Leverage “**white space**” (unused capacity) in existing systems – or **dedicated** stand-alone systems, e.g., **LinuxONE**
- Leverage existing storage systems (**CKD or FB**) – no dedicated storage system required
- Leverage existing IBM Z high-availability and disaster recovery solution by integration of the Accelerator with GDPS/Metro



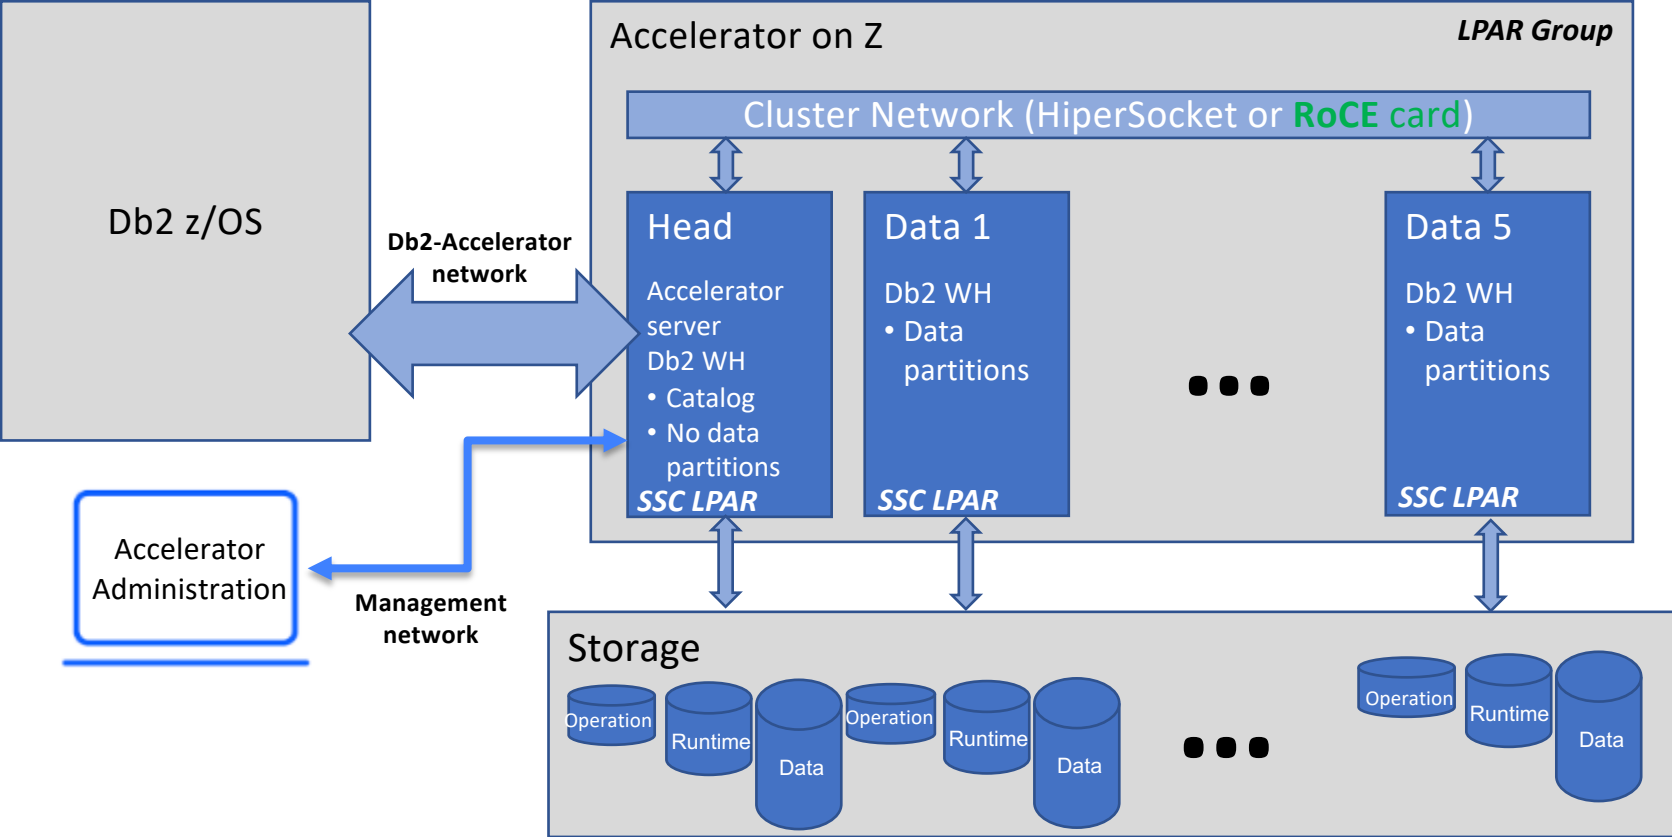
# Single-Node Deployment – Architecture

2 – 35 IFLs (50 IFLs for z16), 64 GB – 4 TB (6,4 TB for z16) memory



# Multi-Node Deployment – Architecture

30 – x IFLs (x=max number of IFLs in system, e.g. 190 in z15, 200 in z16), 1.25 TB – 40 TB memory (total)



# Accelerator on **Z specific** enhancements

*“It is all about performance and stability”*

## Accelerator V7.5.8

- New JSON keywords for better performance: “transient\_storage”, “dispatch\_mode”

## Accelerator V7.5.9

- Support of RoCE cards (RDMA over Converged Ethernet) for the data network and cluster network - *Technical Preview*
- Improved resiliency of the internal network stack
- Improved I/O performance and reduced CPU consumption by exploiting hardware based on-chip compression

## Ability to use local NVMe storage as transient storage (LinuxOne only)

The Db2 Warehouse engine on the Accelerator might write temporary files during query execution if memory is short.

Using high performance storage for these temporary files improves query performance

Local NVMe storage of a LinuxOne system can now be configured to be used as transient storage for temporary files during query execution

- New keyword in “runtime\_environments” in JSON config file: “transient\_storage”:”NVMe”

If specified, all available NVMe devices within a LinuxOne system (LinuxOne local storage) are considered to be used as temporary storage.

- Selected are the ones with an equal block size yielding maximum total size

More info on NVMe devices with LinuxOne:

- <https://www.ibm.com/docs/fi/linux-on-systems?topic=devices-nvme>

## Ability to set dispatch mode for Accelerator (z/OS HiperDispatch)

New optional keyword "`dispatch_mode`" in JSON configuration file.

- Determines how the workload is distributed across the available CPUs.

It refers to the z/OS HiperDispatch mode concept: <https://www.ibm.com/docs/en/zos/2.4.0?topic=mode-concept-hiperdispatch>.

The keyword can be set to "`vertical`" (default) or "`horizontal`".

Vertical dispatch mode means that the workload is concentrated on few CPUs, reducing scheduling overhead and therefore improves performance. This mode is most efficient for systems with many logical processors.

In earlier accelerator maintenance levels always dispatch mode "`horizontal`" was used.

## Accelerator V7.5.9

### Support for RoCE cards for the data and cluster network – *Technical preview*

RDMA over Converged Ethernet (RoCE) is a standard that defines a new RDMA protocol over Ethernet. The RoCE features are PCIe based network adapters on zSystems that provide access to RDMA technology over an Ethernet fabric that provides low latency and high network bandwidth with very efficient host CPU utilization. The RoCE feature can be shared across LPARs.

RoCE cards can be used for [data network](#) as a single card or with a bonding over 2 cards [instead of OSA cards](#)

RoCE cards can be used for the [cluster network of a multi-node accelerator instead of Hipersockets](#)

- Reduces IFL resources of a multi-node accelerator
- Improves stability for internode communication

Further reading:

- <https://www.ibm.com/docs/en/zos/2.5.0?topic=environments-shared-roce-environment>
- <https://www.ibm.com/docs/en/linuxonibm/pdf/l225ro00.pdf>
- <https://www.ibm.com/docs/en/linux-on-systems?topic=networking-roce>

## *Accelerator V7.5.9*

### RoCE cards usage for the cluster network – *Special considerations*

With Accelerator V7.5.9, RoCE cards are supported as technical preview. It is recommended that customers already start using RoCE cards with V7.5.9 for **new Accelerator on Z multi-node installations**. As a rule of thumb start with 2 bonded 25Gb RoCE cards per multi-node accelerator.

For **a future Accelerator release**, the RoCE card usage is going to get a mandatory requirement for new Accelerator on Z multi-node installations.

For **already existing Accelerator on Z multi-node installations** that are upgraded to V7.5.9 or later it is recommended, but not required to use a RoCE card. In case additional IFL resources would be required because of increased workload or data growth, consider to add RoCE cards as a first option before adding IFLs since RoCE cards will replace HiperSocket and free up IFL cycles consumed by HiperSocket

## Accelerator V7.5.9

# Improved I/O performance and reduced CPU consumption

The Accelerator now exploits the ['Accelerated Data Compression'](#) feature available with Linux on IBM zSystems and LinuxONE

- Hardware based on-chip compression

- Results in ***improved I/O performance and reduced CPU consumption***

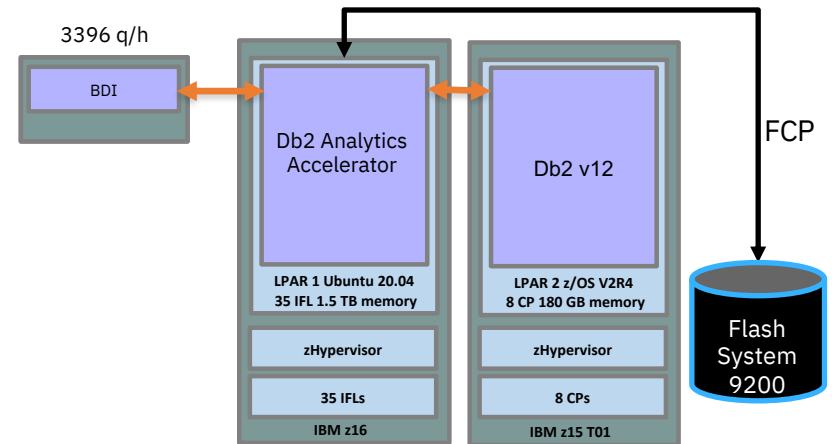
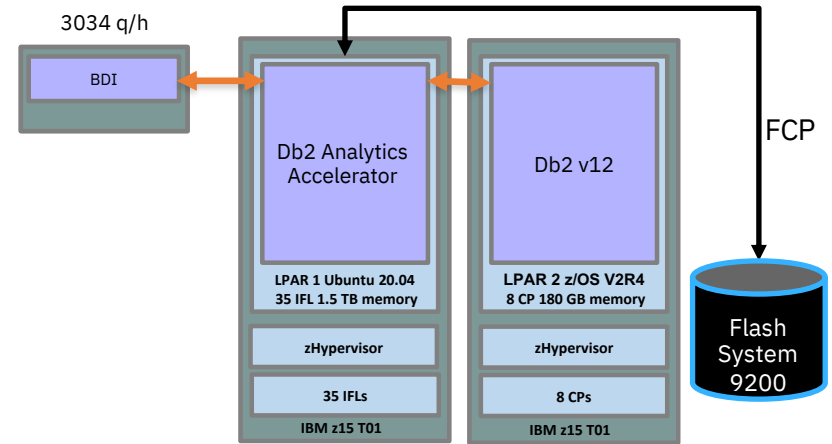
With V7.5.9 exploited for the Accelerator Server process only for internal cataloging and query result set spill-to-disk activities

With V7.5.9 not yet exploited for the Db2 Warehouse component

# Db2 Analytics Accelerator for z/OS on Z on IBM z16 versus IBM z15 T01

Run a workload executing a broad set of OLAP queries with varying complexity on IBM Db2 Analytics Accelerator for z/OS on IBM Z with **11% more throughput** on 35 IFLs on IBM z16 versus IBM z15 T01

**DISCLAIMER:** Performance results based on IBM internal tests running a multi-query analytics workload (fixed ratio of 70% simple, 25% intermediate, 5% complex analytic queries) against the IBM Db2 Analytics Accelerator for z/OS on IBM Z 7.5.7 on an IBM z16 versus on an IBM z15. IBM z16 configuration: Ubuntu 20.04 in an LPAR with 35 dedicated IFLs, 1512 GB memory, and IBM FlashSystem 9200 storage. IBM z15 T01 configuration: Ubuntu 20.04 in an LPAR with 35 dedicated IFLs, 1512 GB memory, and IBM FlashSystem 9200 storage. Results may vary.



# Product Roadmap – IDAA

3Q 2022	4Q 2022	2023
<p data-bbox="40 396 498 448">Focus on performance, reliability, stability, diagnoseability</p> <ul data-bbox="40 494 498 969" style="list-style-type: none"><li data-bbox="40 494 498 579">• <b>Schema change support for non-replicated tables</b> (currently in Tech Preview)</li><li data-bbox="40 620 498 672">• Improvements for Integrated Synchronization serviceability</li><li data-bbox="40 713 498 743">• Query performance improvements</li><li data-bbox="40 784 498 841">• <b>Indication when tables are not used</b> for smarter table management</li><li data-bbox="40 882 498 969">• Improved resiliency and continuous availability with <b>synchronizing tables between Accelerators</b> (Tech Preview)</li></ul>	<p data-bbox="571 396 1045 478">Enabling additional use cases by supporting <b>large object data types (LOB)</b> (staged approach)</p> <ul data-bbox="571 521 1045 860" style="list-style-type: none"><li data-bbox="571 521 1045 603">• Improvements on various monitoring areas e.g., AoT, Performance, Integrated Synchronization, and others</li><li data-bbox="571 647 1045 734">• <b>Customer-accessible way to define, migrate, and change Db2 configuration settings</b></li><li data-bbox="571 778 1045 860">• New use cases for UI for Accelerator administration based on <b>Db2 Admin Foundation</b></li></ul>	<p data-bbox="1093 396 1760 448">Exploitation of Db2 13 for z/OS capabilities as well as IBM z16 hardware features</p> <ul data-bbox="1093 494 1818 647" style="list-style-type: none"><li data-bbox="1093 494 1818 547">• Continue to enhance and support across all deployment options based on customer feedback and Aha Ideas.<ul data-bbox="1170 558 1818 647" style="list-style-type: none"><li data-bbox="1170 558 1818 610">• Exploit HW-acceleration for compression, cryptographic operations and sorting</li><li data-bbox="1170 621 1818 647">• Exploit locally attached NVMe</li></ul></li></ul>

# More Info on Fabric and IBM Z Data Modernization

IBM [Data Fabric entry point](#)

IBM Definition of a Data Fabric: [What is a Data Fabric?](#)

IBM Z application modernization patterns, [Data-centric integration](#).

## Blog Series on Data Fabric and Z

1. [Data Fabric and Z Series - Protecting sensitive data while providing self-service on Db2 for z/OS](#)
2. [Data Fabric and Z series - Monitoring and improving the quality of data coming from Db2 for z/OS](#)
3. [Data Fabric and Z Series - Fuel your Machine Learning models with data on zSystems](#)
4. [Data Fabric and Z Series - Unlock zSystems data with Cloud Pak for Data as Service \(1\)](#)
5. [Cloud Pak for Data 4.5—The Data Fabric platform for zSystems data sources](#)
6. [IBM Db2 for z/OS Data Gate – What is Db2 Data Gate?](#)
7. [IBM Db2 for z/OS Data Gate – Accelerate your queries with Data Gate 2.1](#)
8. [IBM Db2 for z/OS Data Gate –IBM Watson Knowledge Catalog integration](#)
9. [IBM Db2 for z/OS Data Gate – Archiving of Db2 for z/OS data on Cloud Pak for Data](#)
10. [IBM Db2 for z/OS Data Gate – The fire-new installation way for DB2 Data Gate in Cloud Pak for Data 4.5.0](#)
11. [IBM Db2 for z/OS Data Gate – New UI experience of IBM Db2 for z/OS Data Gate 2.1](#)
12. [IBM Db2 for z/OS Data Gate – How to never ever deal with read-after-write inconsistency in replicated databases again](#)

# Db2 Analytics Accelerator V7.5

- Announcement letter for Db2 Analytics Accelerator V7.5
  - [https://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep\\_ca/9/877/ENUSZP19-0299/index.html&request\\_locale=en](https://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/9/877/ENUSZP19-0299/index.html&request_locale=en)
  - <https://www.ibm.com/support/pages/whats-new-ibm-db2-analytics-accelerator-zos-v7>
  
- Release Notes V7.5.1
  - <https://www.ibm.com/support/pages/node/6098626>
- Release Notes V7.5.2
  - <https://www.ibm.com/support/pages/node/6217323>
- Release Notes V7.5.3
  - <https://www.ibm.com/support/pages/node/6336837>
- Release Notes V7.5.4
  - <https://www.ibm.com/support/pages/node/6372182>
- Release Notes V7.5.5
  - <https://www.ibm.com/support/pages/node/6415581>

## Release Notes V7.5.6

<https://www.ibm.com/support/pages/node/6453663>

## Release Notes V7.5.7

<https://www.ibm.com/support/pages/node/6483347>

## Release Notes V7.5.8

<https://www.ibm.com/support/pages/node/6571573>

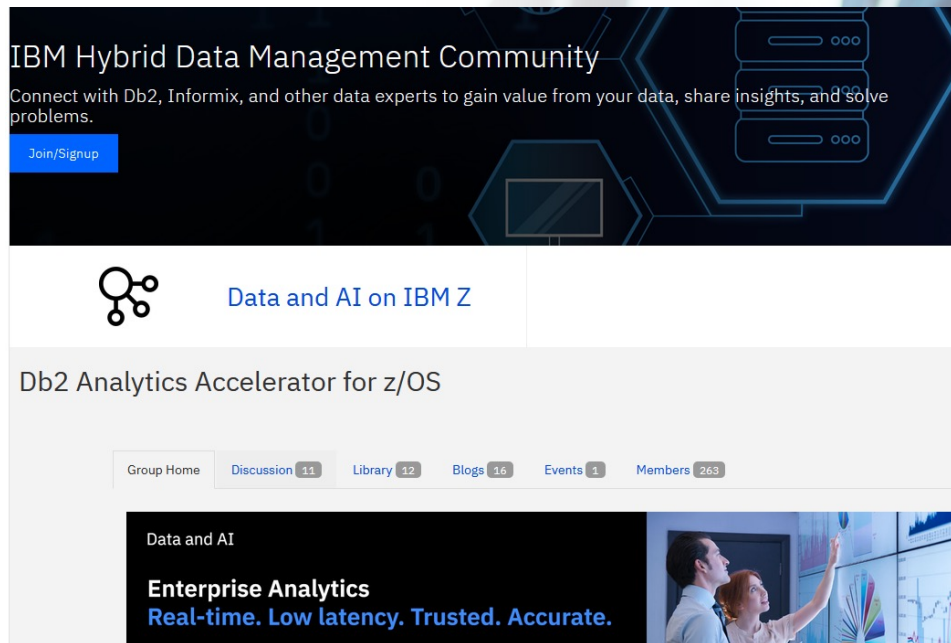
## Release Notes V7.5.9

<http://www.ibm.com/support/pages/node/6607105>

# Community for IBM Db2 Analytics Accelerator

Join Db2 Analytics Accelerator community via


<https://community.ibm.com/community/user/hybriddatamanagement/communities/data-ai-ibmz-home>



IBM Hybrid Data Management Community

Connect with Db2, Informix, and other data experts to gain value from your data, share insights, and solve problems.

Join/Signup

 Data and AI on IBM Z

Db2 Analytics Accelerator for z/OS

Group Home Discussion 11 Library 12 Blogs 16 Events 1 Members 263

Data and AI

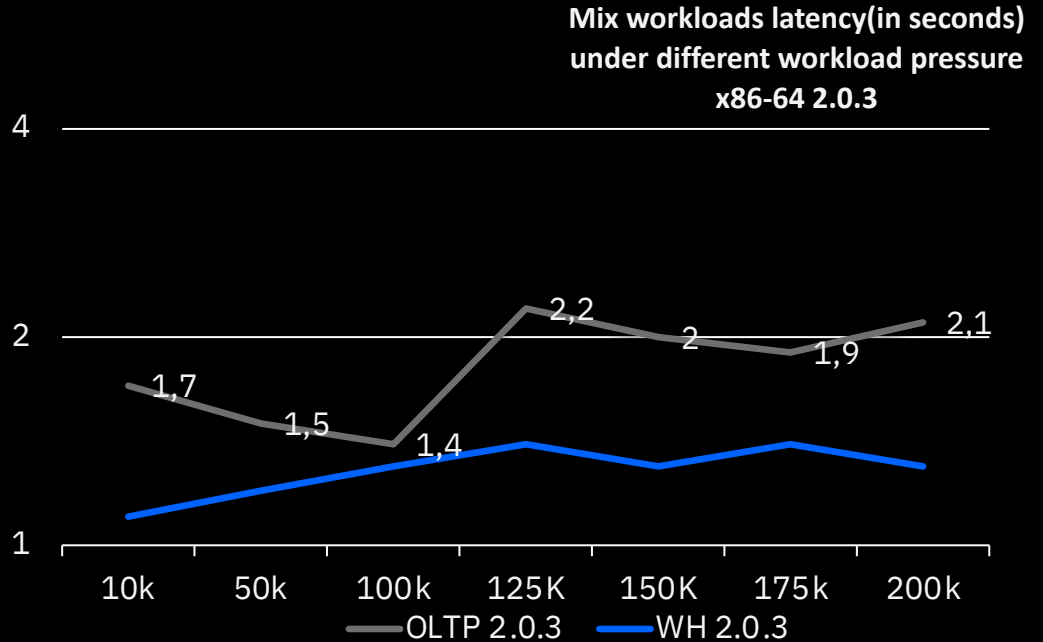
**Enterprise Analytics**  
Real-time. Low latency. Trusted. Accurate.



# Performance – load and synchronization

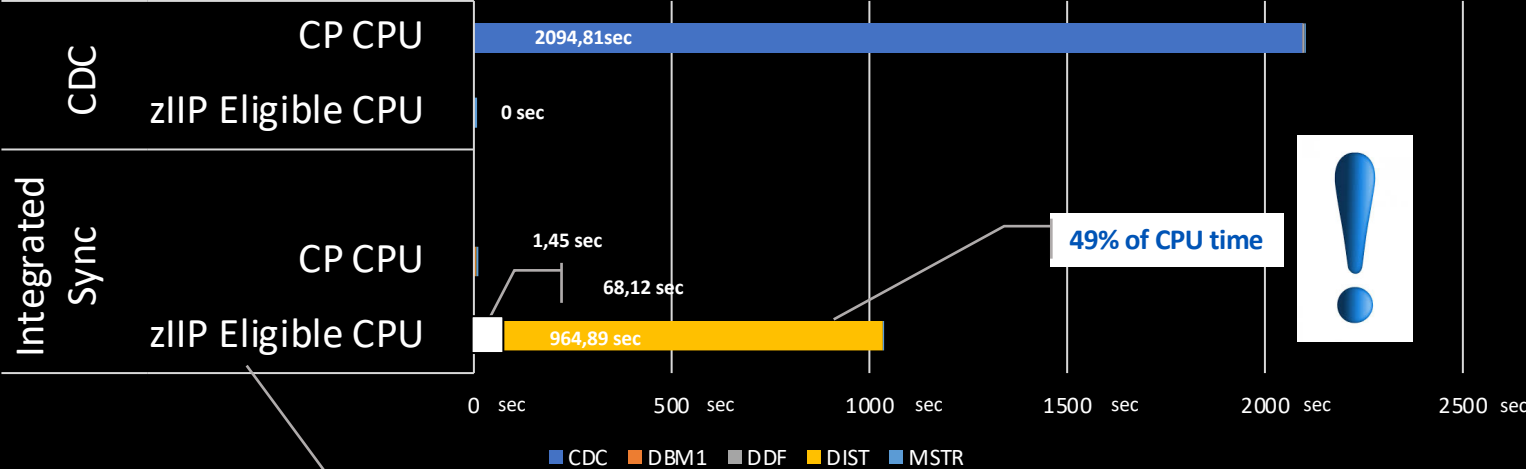
Load – 2TB per hour

Synchronization – Up to 200K IUDs per second




# CP and zIIP Eligible% comparison of Integrated Sync and CDC

CPU Usage comparison between Integrated Sync and CDC at Log read side



Max zIIP eligibility in DIST/DDF address space is 96%



Distributed Data Facility (DDF)  
Distributed Address Space (DIST)