

# Db2 13 for z/OS: Latest Updates

Charlotte zCouncil  
December 7, 2023

Mark Rader


[mrader@us.ibm.com](mailto:mrader@us.ibm.com)

Db2 for z/OS


IBM Z Washington Systems Center



# Foundational and enduring principles: Db2 for z/OS



Exploit deep synergy with hardware and operating system



Drive function down into database layer to simplify application development, reduce data movement, and improve performance




Cache in memory for performance and minimal processing overhead



Eliminate application impact from any event, planned or unplanned

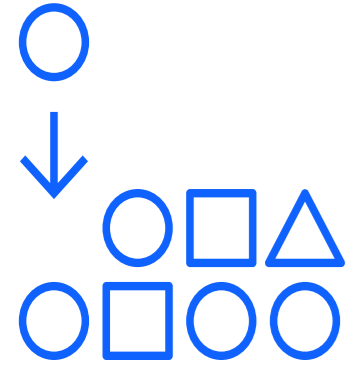
# Db2 for z/OS: technical strategy




Extend foundational strengths in availability, scalability, performance, security, connectivity and resiliency




Extend application and multi-tenancy capability




Expand business insights through in-database AI



Simplify and optimize operational efficiency utilizing AI based automation as needed



Promote and modernize application development and Dev Ops methodology inside and out



Provide secure access to data whenever and wherever it is needed

# Agenda

Db2 13  
timeline

Db2 13  
key benefits

Customer  
experiences

Db2 13  
ecosystem

# Db2 13 function at general availability (GA) – FL 100, FL 500, FL 501

## V13R1M100 (FL100)

- Index look-aside optimization
- Sort optimization
- Expanded SORTL usage with learning from execution (IBM® z15™)
- Improved locking for INSERT to partition-by-growth (PBG) table spaces
- Reduced ECSA storage for IFI buffers
- Reduced agent local below-the-bar (BTB) storage
- DBAT availability improvements
- Improved performance when using external security
- Reduced ECSA storage use for distributed data facility (DDF) processing
- Improved storage monitoring and contraction
- Improved Db2 installation and migration process for customizing the amount of private storage for IRLM locks
- Reduced memory usage for REBIND with APREUSE
- Partition range support in IFCID 306 for users of replication applications
- EDITPROC support in IFCID 306 for users of nonproxy mode replication applications
- Relative page numbering for new PBR table spaces
- Improved default statistics collection granularity
- REPAIR utility WRITELOG for decompression dictionaries
- Enhanced space-level recovery with the RECOVER utility
- Column names longer than 30 bytes
- Db2 support for z/OS continuous compliance
- More concurrent open data sets with z/OS 2.5
- More efficient cleanup for above-the-bar storage
- Index page split instrumentation enhancements
- Accounting information on the longest wait times for common suspension types
- IBM z16 group buffer pool (GBP) residency time
- Subsystem parameter simplification

## V13R1M500 (FL500)

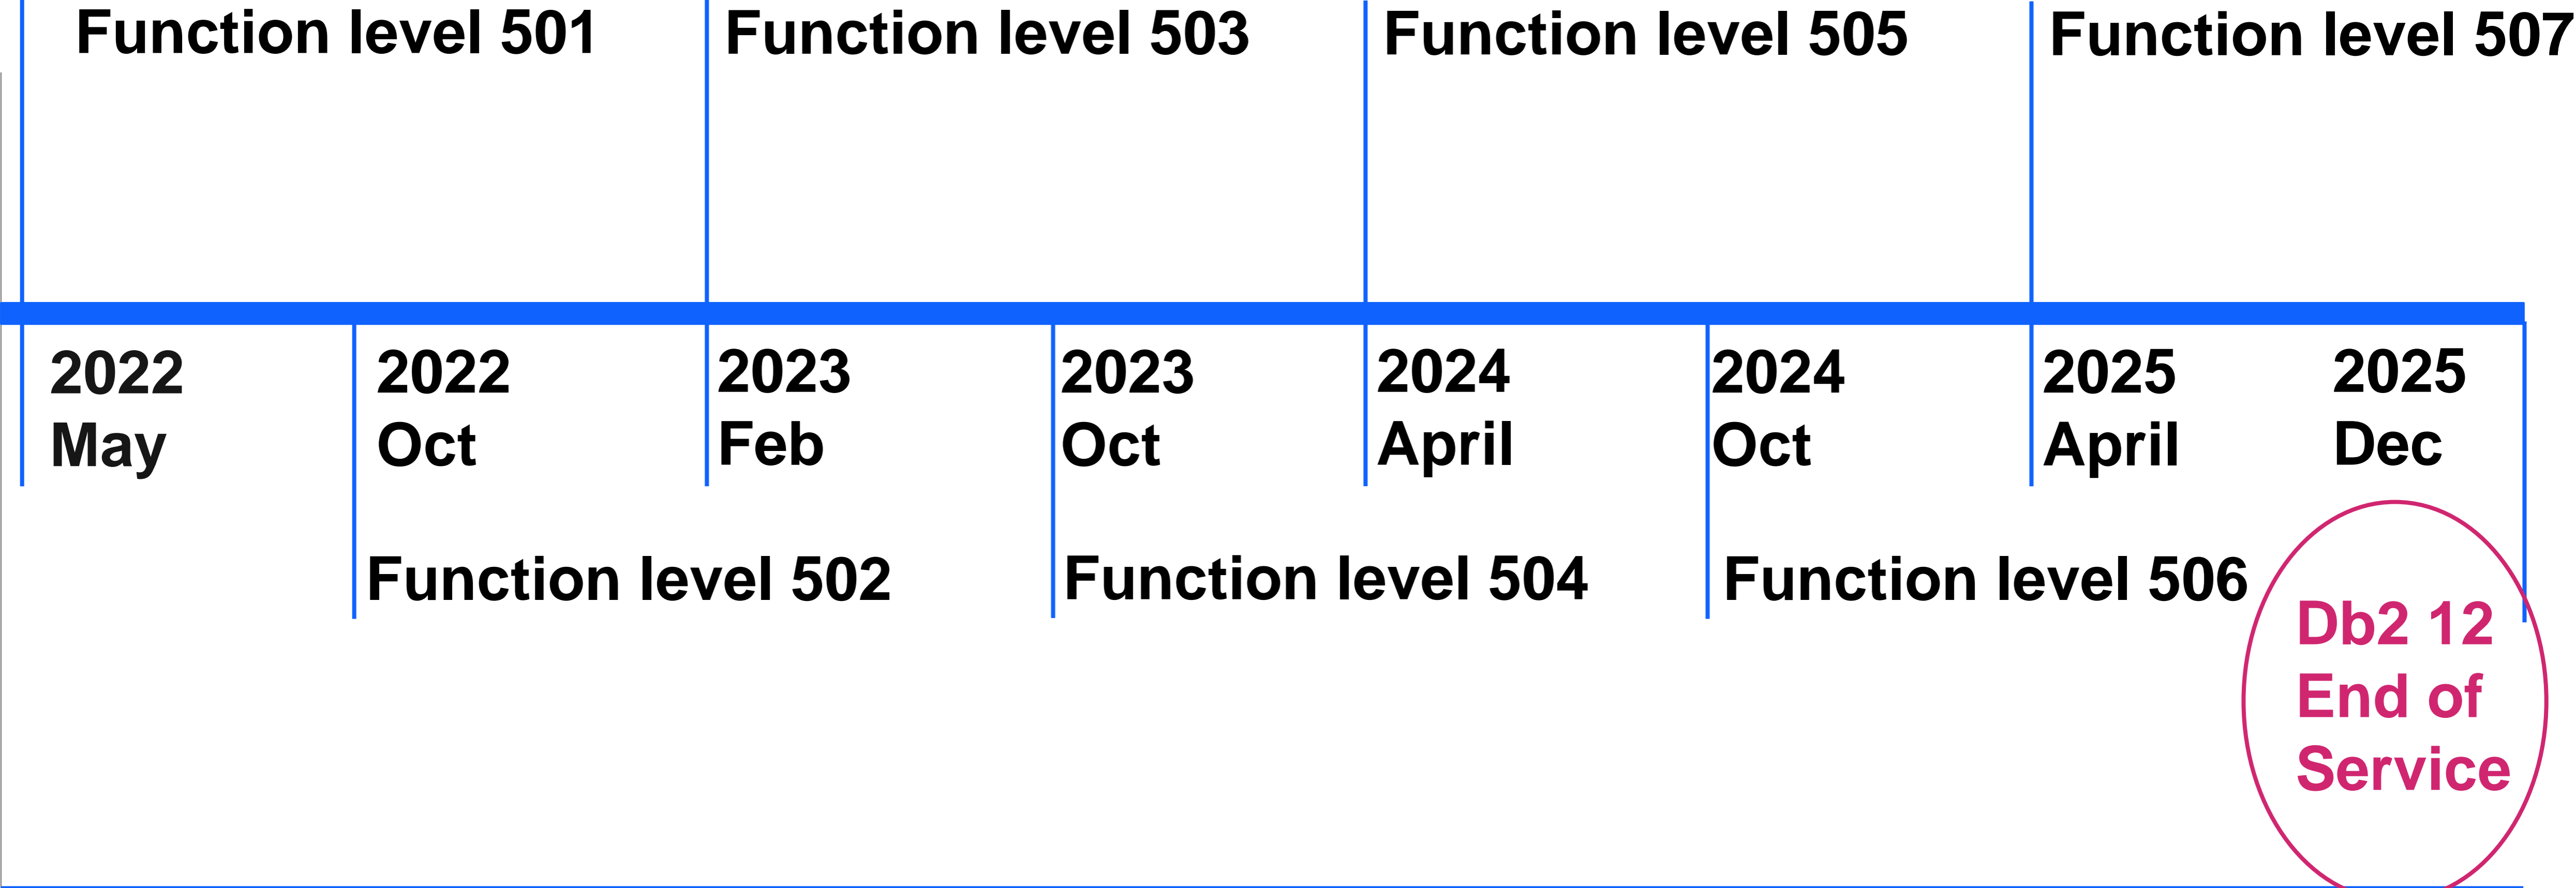
- Increased flexibility for package ownership
- Page sampling for inline statistics
- SQL Data Insights
- Reduced ECSA storage for IFI buffers Part 2
- Online conversion of tables from growth-based (PBG) to range-based (PBR) partitions
- Fast index traversal (FTB) support for larger index keys
- Increased control for applications over how long to wait for a lock
- Profile table enhancements for application environment settings
- Ability to delete an active log data set from the BSDS while Db2 is running
- SPT01 and SYSLGRNX table spaces are converted to DSSIZE 256 GB
- Improved concurrency for altering tables for DATA CAPTURE
- Change REORG INDEX SHRLEVEL REFERENCE or CHANGE so the NOSYSUT1 keyword is the default
- CREATE TABLESPACE uses MAXPARTITIONS 254 by default

## V13R1M501 (FL501)

- Allow applications to specify a deadlock resolution priority
- Profile table enhancements for application environment settings
- Real-time statistics scalability
- Collection of real-time and historical information about utility execution
- Real-time statistics support for index splits

## Db2 13 for z/OS Technology Workshop: Virtual Dec 13-14

# Building upon Db2 13



**Features not dependent upon function level**

**2022**
**2023**
**2024 + . . . .**

# Building upon Db2 13

<p><b>Function level 501</b></p> <ul style="list-style-type: none"> <li>Application based deadlock and timeout</li> <li>Utility execution history</li> <li>Real-time statistics scalability</li> </ul>	<p><b>Function level 503</b></p> <ul style="list-style-type: none"> <li>Enhancement to ROW CHANGE TIMESTAMP</li> <li>Accelerator support for &gt; 32 K IN list</li> </ul>	<p><b>Function level 505</b></p>	<p><b>Function level 507</b></p>
<p><b>2022 May</b></p>	<p><b>2022 Oct</b></p> <p><b>2023 Feb</b></p> <p><b>Function level 502</b></p> <ul style="list-style-type: none"> <li>Statement level invalidation phase 1</li> <li>Db2 controlled Sysplex Workload balancing (Sysplex WLB)</li> </ul>	<p><b>2023 Oct</b></p> <p><b>2024 April</b></p> <p><b>Function level 504</b></p> <ul style="list-style-type: none"> <li>RACF identity token (IDT) support</li> <li>Utility history object level information</li> <li>Trace scalability</li> <li>Statement level invalidation phase 2</li> <li>New SQL Data Insights function</li> </ul>	<p><b>2024 Oct</b></p> <p><b>2024 April</b></p> <p><b>2025 Dec</b></p> <p><b>Function level 506</b></p> <div style="border: 2px solid pink; border-radius: 50%; padding: 10px; display: inline-block; color: pink; font-weight: bold;">             Db2 12 End of Service         </div>
<p><b>Features not dependent upon function level</b></p>	<p><b>2022</b></p> <ul style="list-style-type: none"> <li>System profile monitoring – monitor threads enhancement</li> <li>Db2ZAI V1.5 with granular filtering</li> </ul>	<p><b>2023</b></p> <ul style="list-style-type: none"> <li>Monitor connections for security</li> <li>SQL Data Insights performance</li> <li>Db2ZAI V1.6 simplification</li> </ul>	<p><b>2024</b></p> <ul style="list-style-type: none"> <li>And many more in progress</li> </ul>

# Db2 13 key benefits

## Simpler migration

Smooth process without application impact

## Better availability

Easier schema change, infrastructure management, application change

## Improved performance

Reduced CPU consumption, shorter elapsed time

## Vertical scalability

Increased throughput, fewer constraints

## Enhanced usability

Improved multi-tenancy, simplified administration

## Greater insight

More visibility into important Db2 resources



# Simpler migration



## Continuous delivery

Migration from  
V12R1M510

Continuous delivery  
process mostly  
unchanged

## V13R1M100

CATMAINT required;  
does not change the  
catalog

Optimizer, infrastructure  
features active

Fallback to Db2 12  
possible

## V13R1M500

Most Db2 13 features  
–APPLCOMPAT may be  
required

No fallback to Db2 12  
–V13R1M100\*

Allows CATMAINT  
V13R1M501

## V13R1M501

CATMAINT to  
V13R1M501 required

Complete Db2 13 GA  
features

# Better availability



## Online conversion from PBG to PBR

Take advantage of PBR size, performance and parallelism

One of favorite features for sponsor users and early adopters

## Online removal of active log data sets

For log coverage, encryption, storage migration

No requirement to stop Db2 to run DSNJU003

## Improved serialization for ALTER TABLE

CHANGE DATA CAPTURE no longer quiesces static SQL nor quiesces and invalidates dynamic SQL

## Statement-level dependency and REBIND

Phased approach to invalidate based on statement(s) not on entire package

# Improved performance

## Better INSERT performance

PBG space search

Increased chance to find space or add partition

## Reduced RACF contention

Plan authorization

Global authorization cache

## Increased index fast traversal eligibility (FTB)

Unique index:  
128-byte key

Non-unique index:  
120-byte key

## Greater IFCID 306 log read efficiency

Compression

EDITPROC

## Improved GBP castout processing

GBPOOLT checking

GBP write retry

## Faster REORG INDEX

NOSYSUT1 enforced

Significantly reduced elapsed time and CPU time

# Vertical scalability

## More open data sets

z/OS 2.5 improvement and SWBSTORAGE(ATB)

DSMAX up to 400K

Improved OPEN / CLOSE

## Reduced virtual storage

Instrumentation facility

Distributed: DIST and DBM1

Dynamic SQL with nested procedures, triggers, or UDFs

## Improved storage contraction

Contraction based on monitoring storage thresholds

Less frequent DBAT termination

## Better real storage management

Parameter REALSTORAGE\_MANAGEMENT removed

Db2 monitors free frames

## Quadruple DSSIZE: SPT01, SYSLGRNX

V12R1M500 and REORG

Relieves pressure of more packages and more data sets

## Real time statistics (RTS)

INTEGER to BIGINT

LOCKMAX(0)

# Enhanced usability

## Application control of lock timeouts and deadlocks

CURRENT LOCK  
TIMEOUT special register

DEADLOCK\_  
RESOLUTION\_PRIORITY  
global variable

## System profile monitoring extensions

Current lock special register

Deadlock global variable

RELEASE\_PACKAGE  
keyword

Local thread support

## Subsystem parameters (DSNZPARM)

Some removed

Some default values changed

A few new parameters

## Package ownership flexibility

AUTHID or ROLE based

## Better RECOVER for part-level copies

RECOVER TABLESPACE  
with part-level copies  
without LISTDEF required

## Utility history

Know what utilities ran when; relate to SYSCOPY

Phased approach

FL504 extends to object level

# Greater insight



## Tracking for GBP residency

Directory entries

Data elements

Ability to balance coupling facility(CF) resources

## Tracking for index page splits

Less expensive trace

Information on long duration page splits

RTS updates in FL501

## Tracking lock wait or latch wait

More information on problem during first analysis

Reduced need to recreate the problem

## SQL Data Insights

Built-in functions (BIFs) allow Db2 for z/OS application SQL to use relationships in the data discovered by AI models

# SQL Data Insights

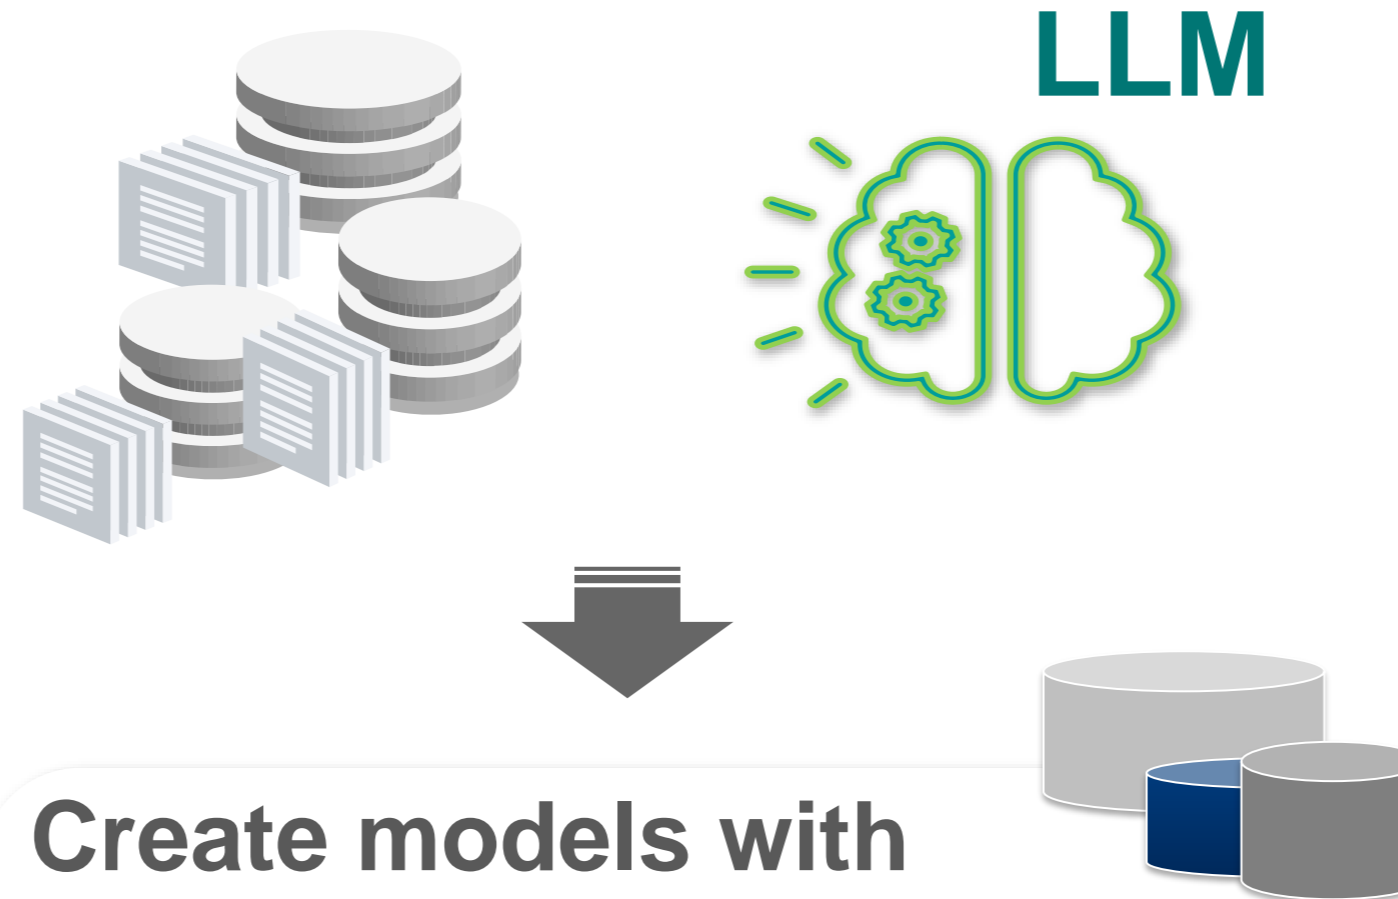
Uncover and monetize hidden insights in Db2 tables utilizing SQL to retrieve vectors generated from Large Language Models

Faster & deeper analysis and action from mainframe data

- Is there similar fraudulent behavior ?
- What is the expected behavior of the new contract?



Rediscover Db2 data



Create models with **meaning vectors** to express relationship and behaviors between rows and columns

AI semantic queries

*Find customers similar to "1234" and rank them*

```
SELECT
  AI_SIMILARITY
  (CUSTOMER_ID, '1234') AS
  SIMILARITY, T.*
FROM CUSTOMER T
ORDER BY SIMILARITY DESC ;
```

**Db2 13 customer experience: favorites**

**Ease of migration**

BMW  
Commerzbank  
Handelsbanken  
ITERGO

**Online REORG  
UTS PBG to PBR**

Commerzbank  
HUK-COBURG  
Volkswagen

**Availability**

–Timeout and deadlock  
ITERGO  
–Online delete active log  
Commerzbank

**Scalability**

BMW  
Commerzbank

**Performance**

–FTB  
ITERGO  
–Real storage  
Commerzbank

**Utility history**

Commerzbank  
ITERGO

**SQL Data Insights**

Citi  
Triton Consulting



# BMW Group

December 2022:  
Db2 13 in  
production in all  
environments

Migration:  
no outage to  
applications

‘Online migration’ for all  
environments

## Utility history

Larger volumes of  
data and  
transactions:  
same MSUs

Below the bar (BTB),  
including ECSA, virtual  
storage relief

## Webcast: early 2023

Thorough discussion of  
environment, processes,  
challenges and  
successes

– [BMW Webcast](#)

# Commerzbank AG

Complete Db2 13  
migration in  
production

Regression testing

Thorough testing in  
multiple environments

frequent communication  
with Db2 development  
in SVL

DSSIZE for SPT01  
and SYSLGRNX

Sponsor user

Disruptive testing

No problems

Online delete of  
active log data sets

Sponsor user

Extensive testing

Feedback for SVL

Utility history

Sponsor user

Already testing next  
phases

MSTR CPU  
reduction due to  
storage management

Initial surprise

Research showed it was  
change from  
REALSTORAGE\_  
MANAGEMENT = AUTO

# Citi and SQL Data Insights

Many years participating in early support and sponsor user programs

Currently in Proof of Concept (PoC)

Promising results

Business users impressed with speed of inference

No data scientist required

# AI with Db2 for z/OS

## Apply AI throughout applications, transactions and operations


### Db2 13 SQL Data Insights (SQL DI)



Uncover and monetize hidden insights in Db2 data

Deep learning Large Language Model (LLM) embedded into Db2, leveraged by tried-and-true SQL queries


### Machine Learning for z/OS



Integrate ML within every transaction -- score in mainframe transactions and batch jobs with performance

Seamless model management - Import, create and deploy the models

### Db2 AI for z/OS (Db2ZAI)



Enhance Db2 database efficiency, resiliency, and performance

Combine machine learning and domain knowledge to enhance Db2 for z/OS operations

# Db2 AI for z/OS

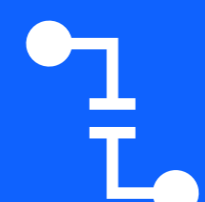
Collect operational data, leverage AI to improve performance and resiliency

Data collection


Db2 for z/OS collects SQL executions, Db2 operational data

Learning and optimization


Db2ZAI learns, predicts and optimizes Db2 execution behavior

 Distributed Connection Control

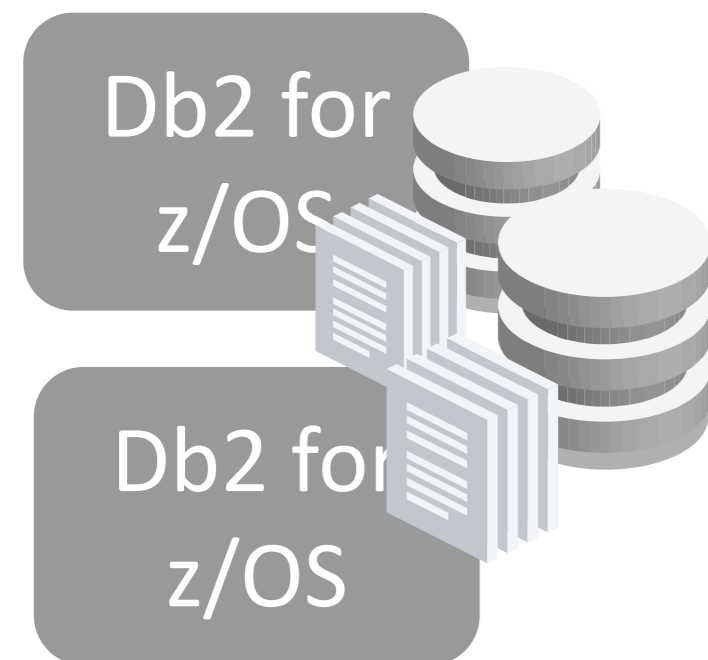
**Protect critical resources**  
*Act quickly, uncover issues early, prevent a flood of connections*

 System Performance Assessment

**Improve Productivity**  
*Less time spent analyzing data, more on finding solutions*

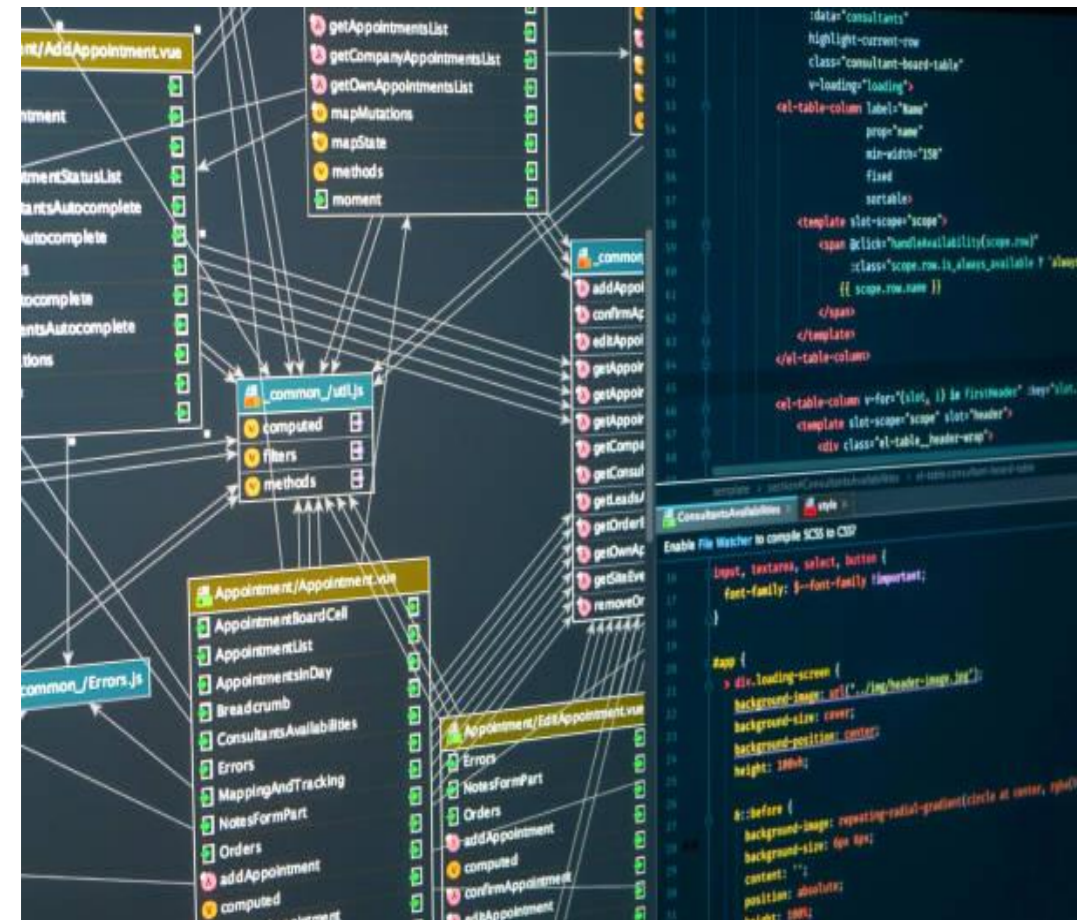
 Enhanced Optimizer Path Selection

**Reduce CPU consumption**  
*CPU saving with better query access path*



# Db2 Administration Foundation & SQL Tuning Services

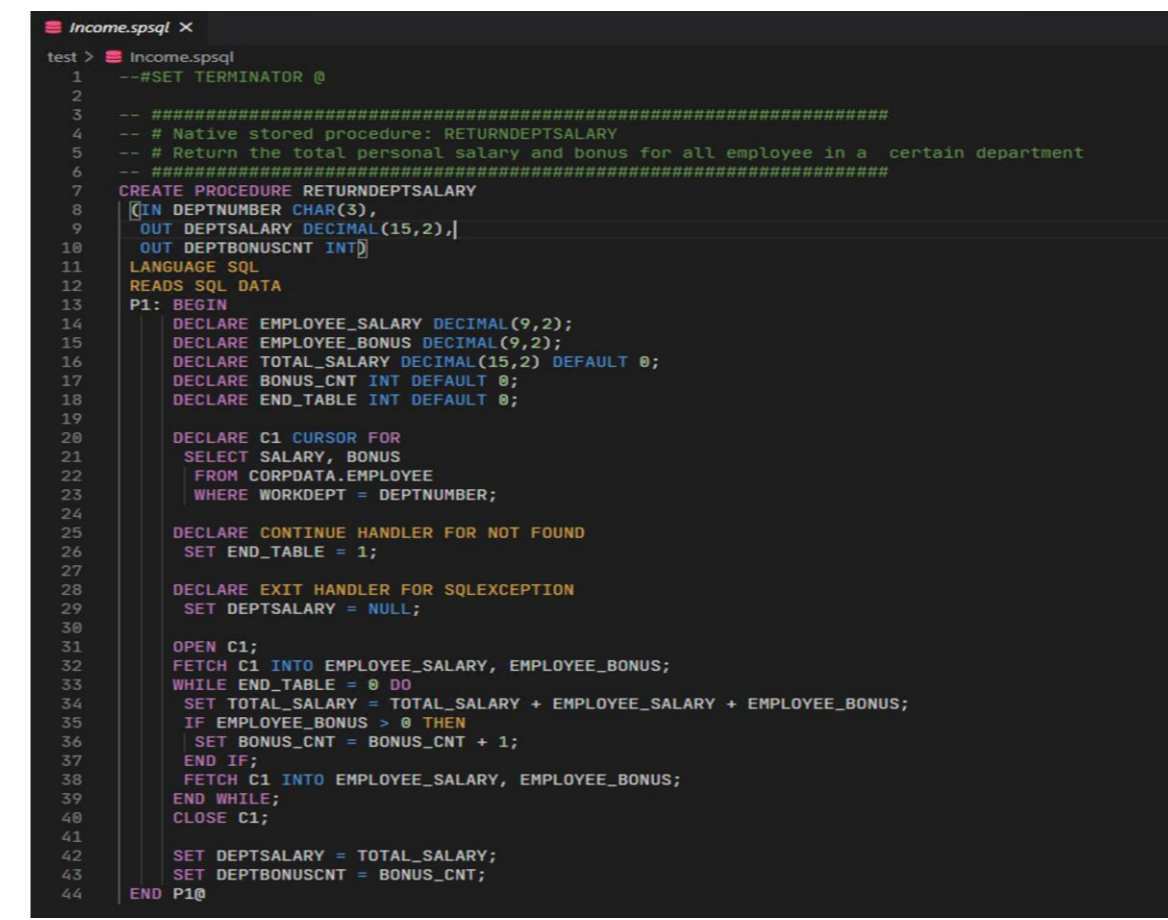
## Discover



## Administer



## Develop & Test



## Tune



- Intuitive search of entire sysplex with visual object representation
- Browser-based user interface built on the open source Zowe Virtual Desktop
- Search across multiple subsystems
- Flexible filtering options & customizable results table

- Dashboard-style object details view
- Db2 Analytics Accelerator administration
- Intelligent Db2 command
- Single-object DDL generation
- Augmented with Zowe: JES, MVS, and USS explorers

- VS Code extension
- SQL Editor with syntax checking reduces time, errors and efforts
- Navigate Db2 catalog
- Execution of SQL with parameter markers and export query results to CSV
- BIND/REBIND with editing

- Single query and workload tuning
- Visual Explain
- Access Path Advisor
- Query rewrite Advisor
- Statement cache + dynamic plan stability features

# Ecosystem spotlight – Analytics Accelerator and Data Gate


★ **Containers on IBM Z**  
e.g., OpenShift, zCX

★ **Data access modernization**  
(REST services, Java, ...)

★ **Db2 Analytics Accelerator**

★ **Consistent data**

★ **Read & write same API**




*IBM Z repositories, relational / non-relational, modern / historic*

**Move the cloud-native workload (containers) to the IBM Z**



**Move IBM Z data to the (private, public, hybrid, ...) cloud for read access**



*Containers and new services with requirement to access data on IBM Z*

★ **Security and governance**  
Copies of data

★ **Data replication**  
Effort, cost

★ **Latency**  
(depending on the technology)

★ **Primarily read**  
How to access IBM Z for update? Consistency?

# A Quick Word on Naming

**Data Gate:** Overarching name for family of products synchronizing Db2 for z/OS data to modern data platforms

**Db2 for z/OS Data Gate:** Z Software offering that enables data replication from *source* Db2 for z/OS database

**Data Gate on Cloud:** PPA Software offering that enables data replication to *target* Db2 Warehouse on Cloud

**Db2 Data Gate Cloud Pak for Data base service:** PPA Software offering that enables data replication to *target* Db2 on Cloud Pak for Data



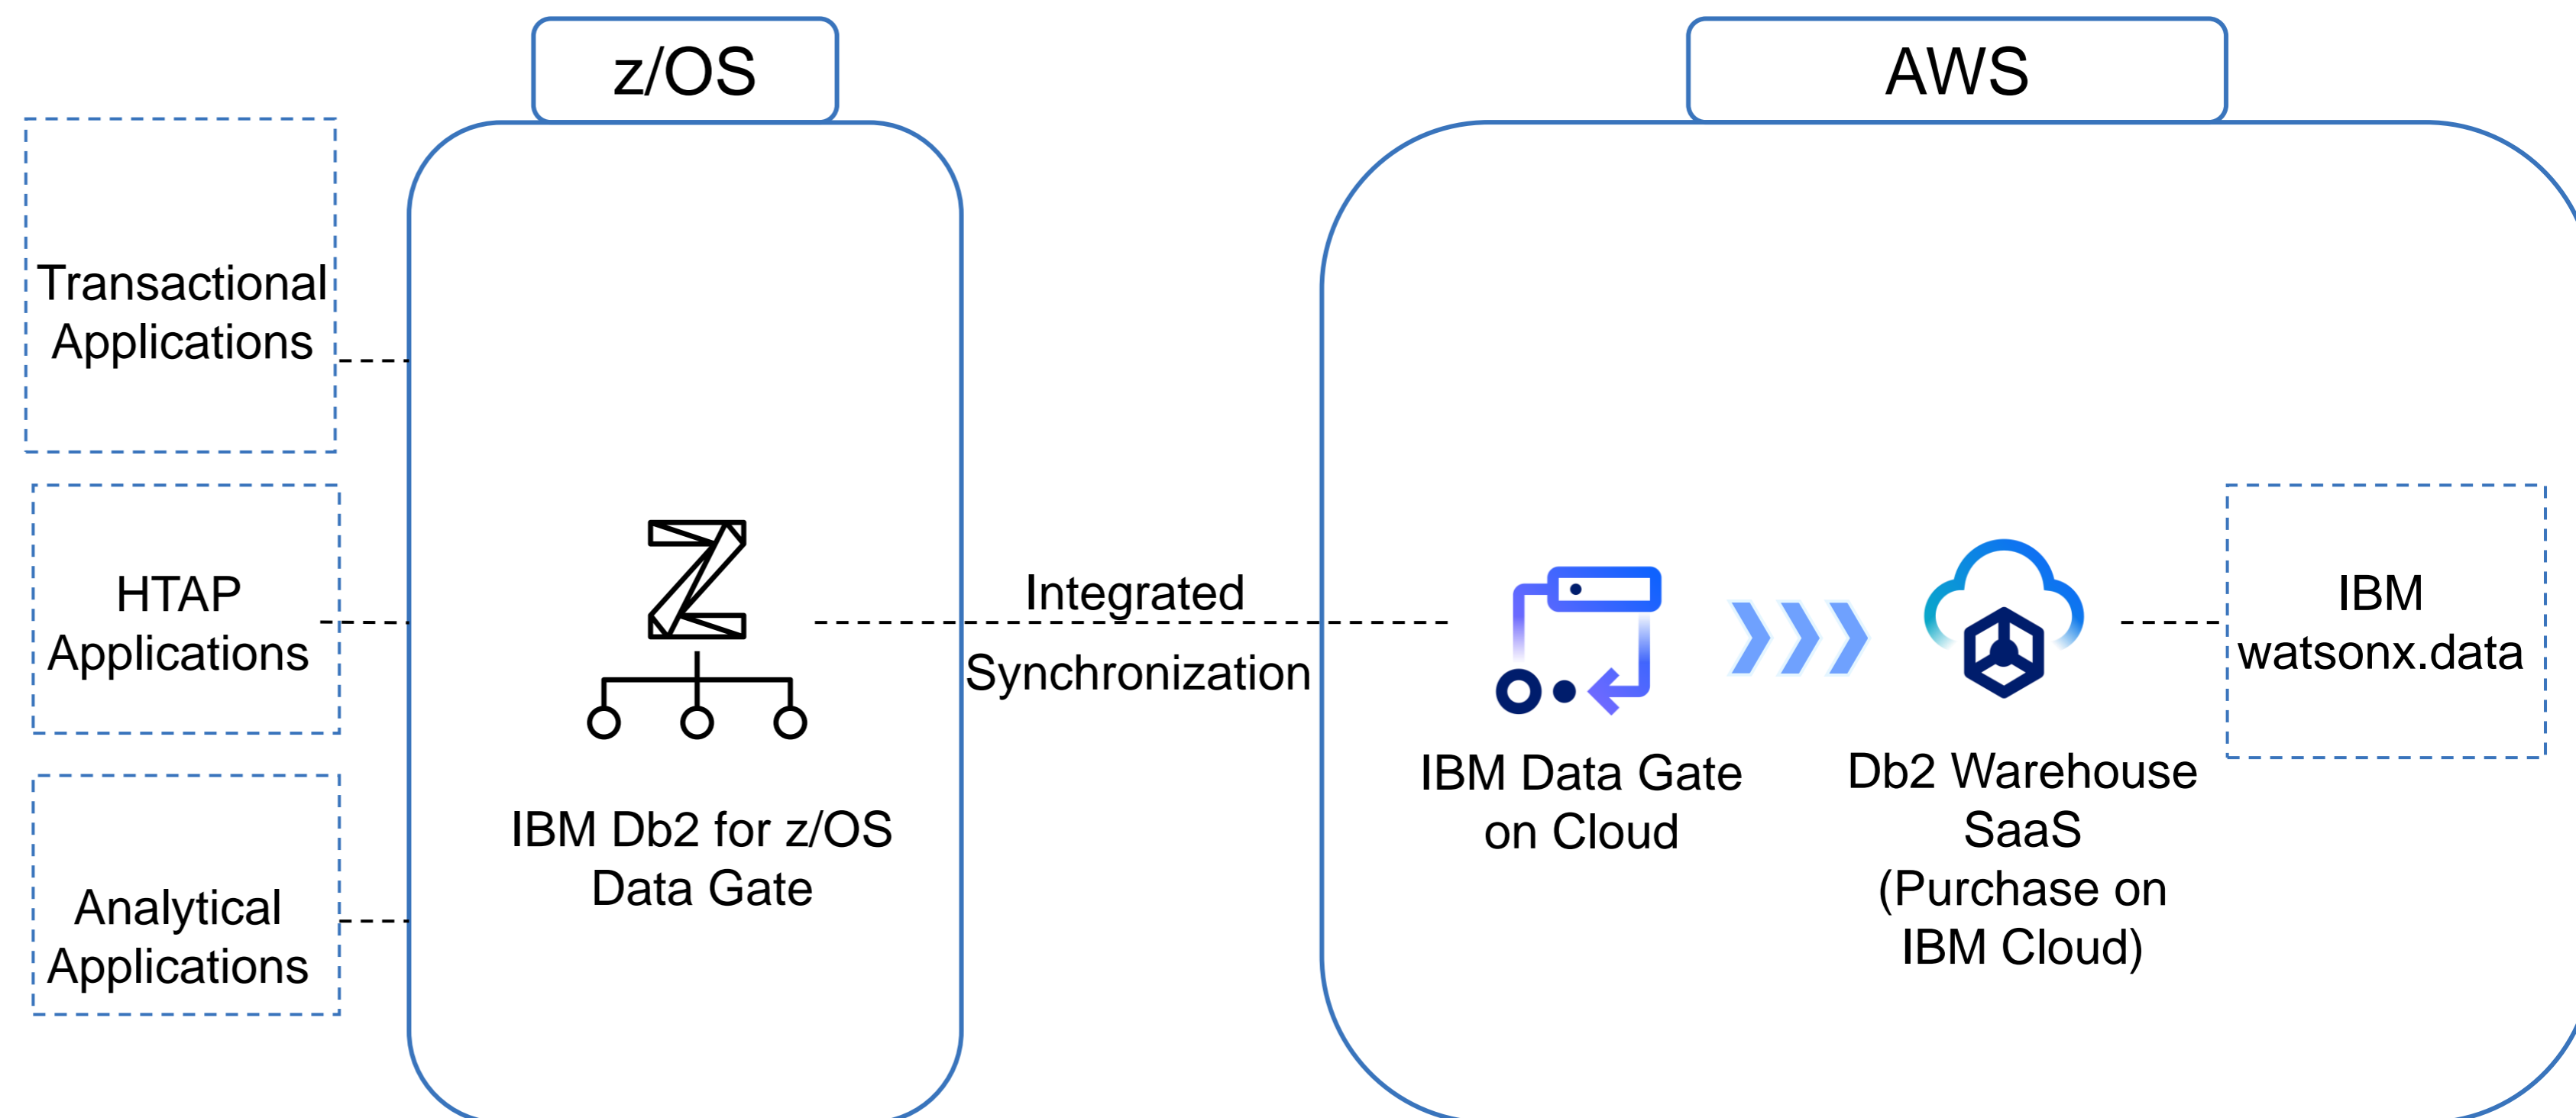


# IBM Data Gate on Cloud

- Synchronize Db2 for z/OS data and make it accessible for:
  - Direct access to Db2 Warehouse SaaS
  - Ingest into IBM watsonx.data lakehouse via Db2 Presto Connectors
- Deployment platform flexibility
- Parallel Z-Part available for Z Sellers

Announce: **October 10<sup>th</sup>, 2023**

GA: **November 24<sup>th</sup>, 2023**



# For more information on IBM products in the Db2 ecosystem...

- **DevOps**

- [IBM Db2 DevOps Experience for z/OS](#)
- [IBM Db2 for z/OS Developer Extension for Visual Studio Code](#)
- [IBM Urban Code Deploy](#)
- [IBM z/OS Connect](#)
- [IBM Data Virtualization Manager](#)
- [IBM Optim Test Data Management](#)

- **Analytics**

- [IBM Db2 Analytics Accelerator](#)
- [IBM QMF](#)
- [IBM watsonx.data](#)
- [IBM Watson Machine Learning for z/OS](#)

- **Data fabric**

- [IBM Cloud Pak for Data](#)
  - [IBM Watson Knowledge Catalog](#)
  - [IBM Db2 Data Gate](#)
- [IBM Data Virtualization Manager](#)
- [IBM Db2 Analytics Accelerator](#)
- [IBM InfoSphere Data Replication](#)

- **IT operations**


- [IBM Db2 Automation Expert for z/OS](#)
- [IBM OMEGAMON for Db2](#)
- [IBM Db2 Administration Tool](#)
- [IBM Db2 Recovery Expert](#)
- [IBM Query Workload Tuner](#)
- [IBM Db2 Query Monitor](#)
- [IBM Db2 AI for z/OS](#)
- [IBM Db2 Administration Foundation for z/OS](#)
- [IBM Z Operations Analytics](#)

- **Security**

- [IBM Security Guardium Data Protection](#)
- [IBM InfoSphere Optim Data Privacy](#)
- [IBM InfoSphere Data Replication](#)
- [IBM InfoSphere DataStage](#)
- [IBM Security zSecure](#)
- [Unified Key Orchestrator for IBM z/OS](#)

# Questions?


# Summary



Db2 for z/OS is the premier data server for availability, scalability, performance, security, connectivity and resiliency



Db2 13 for z/OS extends these characteristics and adds in-database AI for new application insights into existing data



Db2 13 for z/OS is in production at customer sites and many more are migrating to Db2 13



The Db2 for z/OS ecosystem provides a rich set of products to support your applications and processes that interact with Db2 data

# Thank you!

[mrader@us.ibm.com](mailto:mrader@us.ibm.com)

© 2023 International Business Machines Corporation

IBM and the IBM logo are trademarks of IBM 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 [ibm.com/trademark](https://www.ibm.com/trademark).

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.

Client examples are presented as illustrations of how those clients have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

Not all offerings are available in every country in which IBM operates.

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

**IBM**