

# CenCan Db2 User Group

**Spotlight: Db2 z/OS – News from the Lab**

**Jim Pickel, Db2 Development**

**April 12, 2020**

## **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.

# Agenda

- Journey to Db2 12 for z/OS and Beyond
- Introducing Db2 VNext: Key Themes
- Introducing Db2 VNext: Technical Content
- REMINDER: News from the Lab Blog
- Q&A

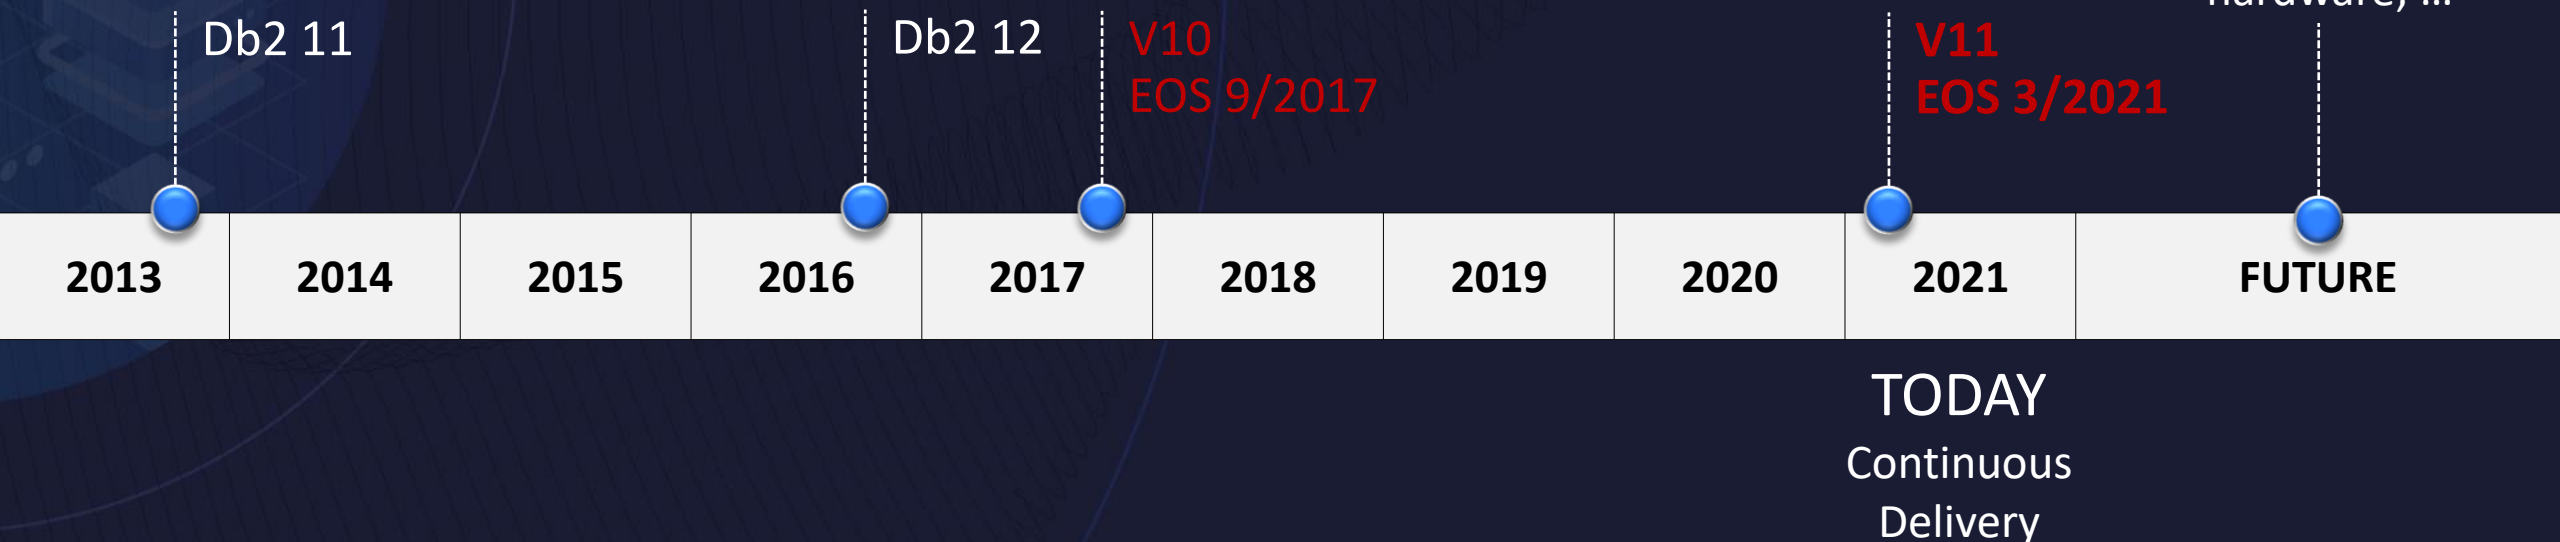
# Db2 for z/OS into the Future

Continuous delivery brings new functions to market 2-4 times per year instead of every 3 years. Db2 12 function level 509 delivered February, 2021.

As of FL509, 9 function levels with 20+ gated features have been delivered, plus over 100 new capabilities since Db2 12 GA

## VNext

When needed due to change in architecture or terms & conditions; retirement of features or hardware; ...



# Db2 12 Continuous Delivery – New Features

## FL 501 – 1<sup>st</sup> post-GA delivery

- LISTAGG

## FL 502 – Apr APAR PI95511

- Transparent Dataset Encryption: Db2 DBA controls
- Casting numeric to GRAPHIC/VARGRAPHIC

## FL 504 – Mar APAR PH07672

- **Huffman data compression**
- New SQL syntax alternatives
- Prevent new deprecated objects
- Passthru of Built-In Functions (OLAP and REGEX) to IDAA

## FL 506 – Nov APAR PH16829

- DROP TABLE automatic drop of explicit table spaces
- Alternative spellings for existing SQL built-in scalar functions to improve compatibility

## FL 508 – Oct APAR PH29392

- **Multi-table table space conversion to PBG**

2017

2018

2019

2020

## FL 503 – Sep APAR PH00506

- **Db2 AI for z/OS (Db2ZAI) – prereqs FL500**
- Migration support on DATA CHANGE OPERATION for temporal auditing
- Enable replication of system-period temporal tables and generated expression columns

## FL 505 – Jun APAR PH09191

- **Rebind Phase-in**
- RUNSTATS sampling simplification
- BIF for Transparent Column Encryption (TCE)
- Temporal and archive on WHEN clause of triggers
- Indexes for Decfloat

## FL 507 – Jun APAR PH24371

- **Application-level locking control**
- Improved statistics management
- CREATE OR REPLACE
- IDAA passthrough expressions

+ Over 100 new Db2 12 features delivered since GA that are not tied to specific Function Level



# Key Themes

- Augmented transactions
- Powered by AI
- Hybrid Cloud

# Augmented transactions and advanced analytics on IBM Z

Inferencing for AI infused insights at the point of client interaction; Advanced analytics next to transaction processing and Systems of Record

## Enhancements enabled through Watson Machine Learning for z/OS

- Auto-feature preparation to accelerate adoption by reducing cost for application development
- Containerized scoring engine that leverage zCX and supports common model types for easy integration with core applications on IBM Z
- Enable large-scale, resource intensive incremental model training through Db2 Data Gate: offline, batch, or online re-training support
- IDAA clients benefit from accelerated AI infused insights derived from complex analytical queries

**No one can scale better than Z**

# Powered by AI

Db2 for z/OS engine with embedded AI capabilities to serve insights and provide deeper integration with Tools

Work with Z Systems team to drive AI Infusion beyond Db2 for z/OS and IDAA

## Db2 for z/OS

- Infuse AI for autonomous database tuning within Db2 for z/OS
- Enable out-of-box access to semantic queries through SQL, leveraging deep learning algorithms from IBM Research that will benefit from on-chip acceleration powered by Watson Machine Learning for z/OS
- Embed machine learning for smarter performance optimization of Db2 online transactions

## Db2 AI for z/OS (Db2ZAI)

- Deliver common set of “curated ready for AI” services that can be exploited to advance and further optimize surrounding ecosystem – e.g., Db2 Z Tools portfolio, IZOA, zSMS, ...
- Provide workload optimizations triggered from system-level events identified by Db2ZAI, powered by WMLz
- Continue to expand security configurations, rapid serviceability, and high availability while improving user experience in response to customer feedback

**No one can scale better than Z**





# Z leadership in hybrid cloud

Integrate data serving technologies across use cases: IDAA, DVM, Db2 Data Gate on the Cloud Pak for Data stack

## Cloud Pak for Data on Z

- Deliver foundational services to support Db2 Data Gate use case in first iteration: Db2 AESE, Db2 WH, Db2 Console
- Continuous delivery of additional services supported on IBM Z/LinuxOne

## Db2 Data Gate/IDAA

- Complete feature gaps to ensure performance and data consistency
- Enable query acceleration use case through Db2 for z/OS - processing on Cloud Pak for Data (Db2 Data Gate)
- Expand targets to include as a service on IBM Cloud, Azure and AWS
- Integrate Db2 Data Gate with Cloud Pak for Data services, e.g., WKC

## Db2 for z/OS and Tools

- Drive advancements in Db2 for z/OS as best-in-class hybrid cloud server with online creation of unique indexes, on-line schema support and re-partitioning
- Transform administration of Db2 through new Db2 Administration Experience

**No one can scale better than Z**

# Why VNext

- Opportunity to make necessary structural changes that cannot fit into Db2 12 maintenance stream
  - Control block changes requiring recompiles, etc
- Opportunity to simplify internal logic
  - Better quality
  - Better reliability
  - Less complex code
  - Improved performance through shorter code paths
  - Simplified development, thereby greater development and service velocity and productivity
- Opportunity to deliver foundational capability that cannot easily be achieved through V12 maintenance stream
- Building upon, not abandoning, Db2's continuous delivery model

# Timeframe

- Db2 VNext GA tentatively planned for early 2022
- Delivers exploitation of new h/w features
- **Guiding principle: Make migration as simple and risk-free as possible**
  - Introduce few to no incompatible changes that may inhibit migration
  - V10, V11, V12 APPLCOMPAT levels continue to be supported in VNext
  - PRIVATE\_PROTOCOL and BIF\_COMPATIBILITY zparms will continue to be supported
- Db2 will ensure adequate time before EOS of Db2 12



# Planning and Migration

- Function levels are consecutive and cumulative
- Db2 VNext will remain consistent with Db2 12 CD model
- Function levels are reset for Apollo and will mirror V12 behavior
  - Migration to V13R1M100, then to V13R1M500, 501, ...
- Ergo, V13R1M100 will contain no new function & be functionally equivalent to V12R1MLAST
- New function development focus shifts to VNext as the base
  - No further function delivered in function levels on Db2 12 after V12R1M509
  - Few new functions delivered in Db2 12 after V12R1M509
    - Minimal exceptions possible
- Migration to V13R1M100 can only be from V12R1MLAST
  - Ensures co-existence & functional consistency
- Co-existence supported, fall back supported



# Primary VNext Candidates

- Simplified Migration
- Smarter, autonomous Db2
- Transformed developer experience
- Always-on improvements



# Simplified Migration

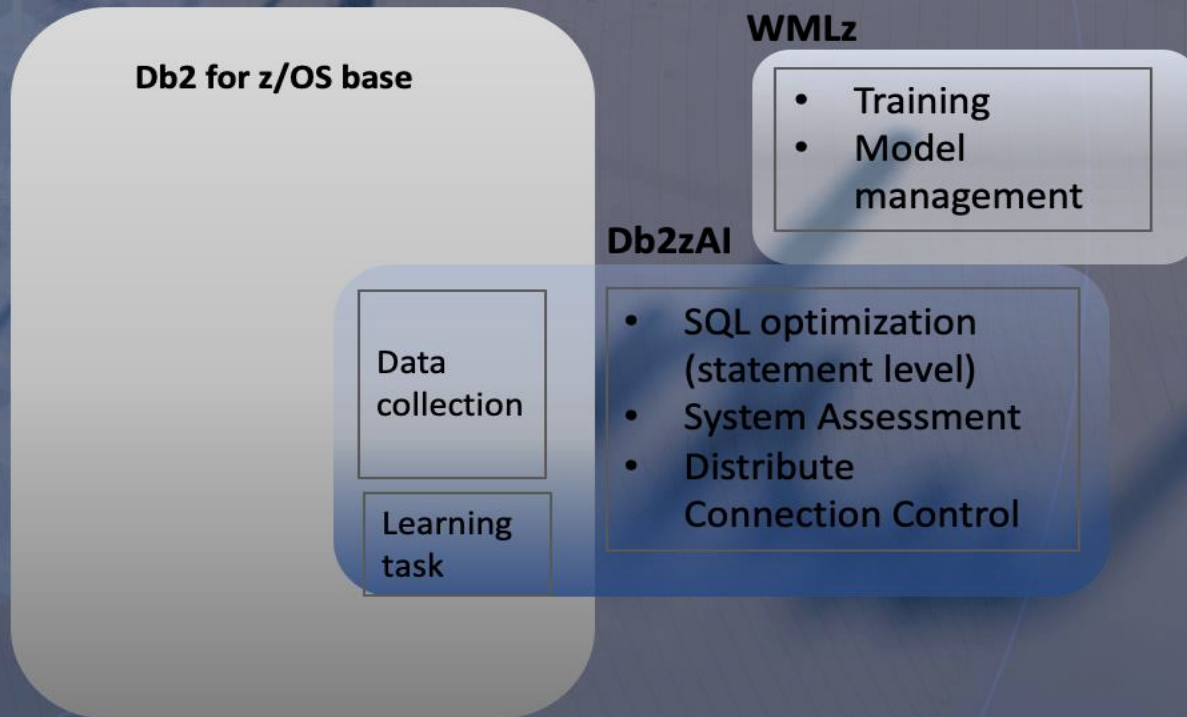
- You indicate which Db2 system is to be migrated
- We discover all configuration values for that specific Db2 system
  - Eliminate need to navigate thru multitude of migration panels
  - Prompt you with values, indicating if any need to be changed for this migration
- We build artifacts to migrate that system
  - Traditional Jobs
  - z/OSMF artifacts
  - Build automation process
    - z/OSMF workflow
    - Other automation method
  - Other artifacts?
- You start the migration, exploiting online migration enhancements
  - Optionally:
    - With user guidance
    - Automated

## Online migration enhancement

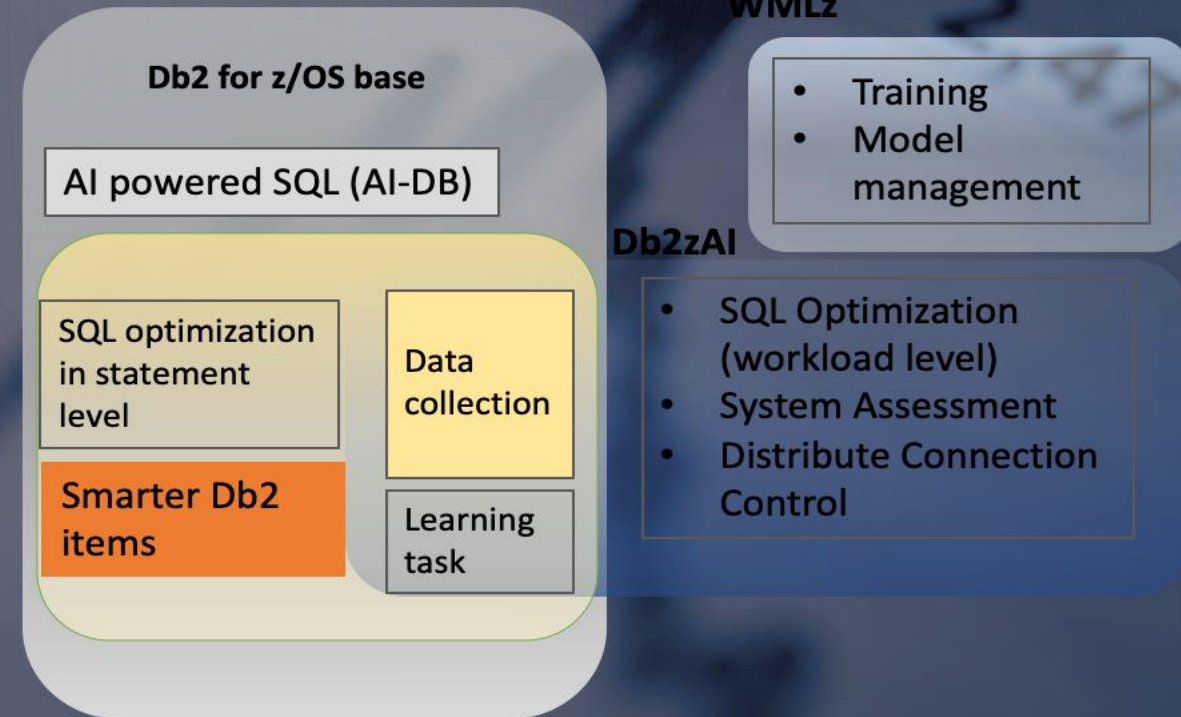
- Display blockers/blocking threads

# Concept of AI embedded Db2z

Today



Vnext





# Smarter Db2 learning from execution data / history

- Smarter index split for insert performance
- Further Access path improvement in statement level
- Smarter Index look aside
- RDS Sort optimization
- Streamline Resource Limit Facility for OLTP queries
- Insert performance on TS by learning from execution

Cost savings and performance improvement through self-tuning



- Open data set management : predictive data set opens
- Simplify workfile management
- Improved IDAA offload decision
- Improved parallelism management
- Shared RID pool optimization
- Reduced requirement for RUNSTATS


Simplification through self-managing





# Transforming the Db2 for z/OS Developer Experience

Function	Strategic direction
Automation	Robust REST APIs to encourage automation, integration with self-service, CI/CD pipelines
Developer	Deliver Db2 for z/OS support via extensions to market- leading cloud native IDE (VS Code, Eclipse Theia)
Administration	No-charge, multi-platform Db2 Management Console. Basic Db2 for z/OS usage, navigation, single query tuning capability
Advanced Administration	Transformed Db2 Tools based on Zowe



# Always On & Other Key Enhancements

- Minimize impact of package invalidations
- DDL break-in with no application impact
- Online conversion from PBG to PBR
- Application-level lock timeout control
- Semantic query support
- Automatic SSL setup
- Address scalability concerns
- Online deletion of logs

# Visit Db2 for z/OS News from the Lab blog

Subscribe for the latest news about enhancements and new capabilities in Db2 for z/OS, from the IBM experts who design, build, test, and support Db2.

How to subscribe and get notified about new blog entries, join the Db2 for z/OS group:

- Go to <https://www.ibm.com/community/> and click Sign in/Join to create your account
- Go to the Db2 for z/OS Group Home page and click Join Group
- Go to the notification settings for your account, find the Db2 for z/OS group in the table, and set your notification preferences
- Find our latest blog entries on the blogs tab or search for the #Db2Znews tag

## Db2 for z/OS Director's Corner: Fast Start to 2020!

By [Mo Townsend](#) posted 2 days ago

5 Recommend



2020 has already been a busy year for us in Db2 development, and we hope it's been productive for you too! We have a lot of news to share, so I thought I'd give you a peek at what's happening, as we set course to focus on our 2020 priorities: Attracting new applications; ensuring the quality that Db2 for z/OS is known for; and refining our product and process as some talent in our team transitions to retirement and new talent takes the helm.

### Db2 Development team transitions

I have two announcements regarding leaders whom many of you know: Jeff Josten and Chris Crone have retired from IBM.

**You can also continue to follow us on Twitter**  
**@IBMDb2ZLabNews**, and you can find our  
latest blog entries at [ibm.biz/db2znews](https://ibm.biz/db2znews)

### IBM Hybrid Data Management Community

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

#IDUGDb2



## Spotlight on PBG-PBR Conversion

- Why?
  - PBG is the most common UTS table space type. Most table spaces converted to UTS are PBG. PBG works well for small/medium sized tables but there are a number of drawbacks once the table grows too large, such as
    - Insert performance degradation
    - Query performance degradation
    - Regaining clustering requires REORG of the entire table space
    - All indexes are NPIs, with resultant issues associated with very large indexes
    - No utility partition parallelism and limited utility part-level support
  - A DBA decision taken in the past to place a table in a PBG table space cannot be rectified without a major application outage
  - Lack of online conversion support can result in increasingly degraded application performance and degraded ability to manage large Db2 tables



## Spotlight on PBG-PBR Conversion

- As-is
  - A DBA determines that it is no longer sustainable to maintain a table in a PBG table space
  - They must either
    - Incur a major production outage to unload the data, drop and recreate the table space and table and associated objects and reload the data
    - Split the data into smaller tables, requiring application changes to cope with a DBMS deficiency
    - Live with the increasing challenges associated with large PBG table spaces in the hope of an IBM solution

# Spotlight on PBG-PBR Conversion

- To-be

1. ALTER TABLE
2. REORG TABLESPACE

```
ALTER TABLE SCR001.TB01 ALTER PARTITIONING TO  
PARTITION BY RANGE (ACCT_NUM)  
( PARTITION 1 ENDING AT (199),  
PARTITION 2 ENDING AT (299),  
PARTITION 3 ENDING AT (399),  
PARTITION 4 ENDING AT (MAXVALUE) );
```


## Spotlight on PBG-PBR Conversion

- Preliminary notes
  - Apollo M500 only
    - Technically APPLCOMPAT M500
  - Conversion to PBR RPN only
  - Enforce MAXVALUE|MINVALUE for last partition
    - Prevent risk of discarded records by REORG
  - No support for tables with XML or LOB columns (initially)
  - No conversion of NPIs to PIs



Q&A





Speaker: Jim Pickel  
Company: IBM