



## NEWSLETTER 05 | AUGUST 2023

### Contents

Virtual Db2 User Group Presentation	1
Next Meeting	3
News and Announcements	4
About the Virtual Db2 User Group	4

### Virtual Db2 User Group Presentation

## ENHANCING DB2 DATABASE DESIGN AND APPLICATION PERFORMANCE: A COMPREHENSIVE OVERVIEW OF DB2 12+ AND 13 FEATURES

In her recent presentation, internationally renowned consultant and Db2 expert, Susan Lawson, took a deep dive into the latest features and improvements introduced in Db2 version 12 and version 13 for z/OS. The presentation aimed to educate database administrators and application professionals on leveraging these enhancements effectively to enhance and boost database design and application performance. Susan's dedication to staying up-to-date on all APARs and function levels ensures that organizations can maximize the benefits offered by Db2.

To kick off the presentation, Susan gave an overview of the updates from Db2 12+/13 as they related to database performance and availability. These included features and changes such as:

**YI&A Database Performance and Availability – Db2 12+/13**

- Relative Page Numbering Default
- Improved Insert for RPN in Data Sharing
- PBG to PBR Conversion
- Remove Stack Limits for Pending Alters
- Max Parts Default Change
- Increased Number of Open Data Sets
- Improved Data Capture
- Larger Column Names
- Non-Deterministic Expression Column Auditing
- Segmented Multi-Table Tablespace to UTS
- UTS PBR Hidden ROWID
- Faster PBG Insert
- Improved Index-Lookaside
- FTB Support for Non-Unique Indexes
- FTB Selective Index Usage
- FTB INDEXTRAVERSECOUNT
- Online Load Replace
- Trigger, Temporal and Archive Transparency
- RTS Scalability Improvements
- LOB Compression Improvements
- Transparent Data Encryption Functions
- Huffman Compression

Db2 12+/13 for z/OS: Database Design and Application Performance Features and Usage 7

**Relative Page Numbering Default:** The introduction of relative page numbering for PBR tablespaces in v.12 removed limitations when building large tablespaces. In v.13, relative page numbering became the default, requiring extended addressability on datasets.

**PBG to PBR Conversion:** Db2 v.13 introduced the ability to convert from PBG to PBR tablespaces (Function Level 500) using a new ALTER PARTITIONING to PARTITION BY clause, allowing better data distribution and sizing opportunities.

**Online Load Replace:** The new Online REORG SHRLEVEL REFERENCE for LOAD Utility in v.12 enables a seamless and transparent switch between tables during a LOAD operation, improving availability during data loading.

**Trigger, Temporal, and Archive Transparency:** Triggers can now reference System Temporal and Archive Enabled Tables, enhancing data tracking and auditing capabilities.

**RTS Scalability Improvements:** Data types in RTS Stats have been enhanced to accommodate larger amounts of data, providing better scalability for statistical data collection.

## SPONSORS





## ENHANCING DB2 DATABASE DESIGN AND APPLICATION PERFORMANCE: A COMPREHENSIVE OVERVIEW OF DB2 12+ AND 13 FEATURES

Continued

**Improved Index Look-Aside** is a crucial performance feature in Db2 that has evolved over time. In versions prior to v.13, it only worked for INSERTs and DELETEs on highly clustered indexes. However, in Db2 v.13, Index Look-Aside is enhanced to work for INSERTs, UPDATEs, and DELETEs, regardless of cluster ratio or statistics accuracy. Db2 dynamically adjusts its usage of Index Look-Aside based on patterns, ensuring optimal performance. This improvement helps reduce maintenance costs associated with introducing indexes, as it mitigates the overhead of INSERTs, UPDATEs, and DELETEs. However, it's essential to be mindful of the number of indexes added to a table to avoid unnecessary overhead.

Another enhancement that Susan emphasized is one of the most significant enhancements in Db2 13—an allowed **Increase in the Number of Open Datasets**. With larger objects, tablespaces, and indexes, the need for open datasets has grown. Db2 13 increased DSMAX, the Dynamic Storage MAXimum, from 200,000 to 400,000. This improvement enables Db2 to handle more concurrent data sets or other activities like concurrent threads, reducing the risk of datasets closing prematurely and causing performance issues.

**Y&A** Increased Number of Open Datasets

- Issue
  - With large increases in data sets needed there is a growing need for a larger number of open datasets
    - Due to business consolidation, Db2 member consolidation, data growth, and conversion of segmented table spaces to PBGs
  - When DSMAX is reached (within 3%) datasets get closed
    - If reopened, can cause performance problems in applications
  - Db2 11(Db2 10 via APAR PM88166) increased the maximum number from 100,000 to 200,000
  - Limit of concurrent open data sets is due to
    - Required 31-bit DBM1 private memory per open data set
    - Amount of 31-bit private memory remaining in address spaces after subtracting required common memory
- Db2 13 and z/OS 2.5
  - z/OS 2.5 moves a portion of the control blocks for open data sets ATB
    - DSMAX increases from 200,000 to 400,000
    - Db2 can use the extra room to increase number of concurrent data sets or other Db2 activities (i.e. concurrent threads)
    - Can improve performance by decreasing the need to constantly open/close datasets

**DSMAX**

Db2 12+/13 for z/OS: Database Design and Application Performance Features and Usage 16

In the latter half of her presentation, Susan discussed various application and SQL enhancements that included:

**Application Lock Timeout Control:** Allows users to control lock timeouts per application, overriding system zparm settings.

**Global Variables for Locking Limits:** This feature offers greater control over locking limits, improving concurrency in Db2.

**Statistics Collection Improvements:** Enhanced statistics collection enables more effective index and data management.

**Reduced Sort Row Length:** Db2 13 optimizes sort processing by reducing the row length.

**Y&A** Deadlock Priority Control

- Prior to Db2 13
  - When different lock requests from multiple threads a deadlock occurs
    - IRLM resolves deadlock by choosing a *victim* whose request is then denied
  - Application cannot influence which thread is chosen as deadlock victim
  - Batch data definition (DDL) jobs frequently fail due to deadlocks
    - Failing the DDL statements instead of other SQL is not optimal
  - Restart or retry logic and careful scheduling can help
    - Restart/retry logic requires code change and additional processing and is often not done
    - Scheduling is increasingly difficult to optimize in continuous availability environments
- Db2 13 (M501)
  - New DEADLOCK\_RESOLUTION\_PRIORITY built-in global variable **DEADLOCK\_RESOLUTION\_PRIORITY**
  - Allows for specification of a deadlock resolution priority value to use in resolving deadlocks with other threads
    - 0-255, NULL, DEFAULT
    - Higher value = less likely lock requests by application will be the victim in a deadlock situation

**SET SYSIBMADM.DEADLOCK\_RESOLUTION\_PRIORITY = 100**

Db2 12+/13 for z/OS: Database Design and Application Performance Features and Usage 55

One notable feature of Db2 13 is **Deadlock Priority Control**. Prior to Db2 13, IRLM would resolve a deadlock by choosing a “victim” whose request is then denied. Thanks to the new built-in global variable called DEADLOCK\_RESOLUTION\_PRIORITY, you can set a priority for when Db2 is trying to determine who the victim is going to be in a deadlock. The higher the priority value, the less likely that application or that job will become the victim in a deadlock.

## SPONSORS



**NEXT VIRTUAL MEETING - September 26, 2023, 10:30 AM CDT**

### Db2 Level Up

Join us for an enlightening event as we explore the intricacies of Db2 Upgrades and shed light on the often encountered phrases such as “Code level,” “Catalog level,” “Function-level,” “Application compatibility level,” and “Connect level.” These terms can often raise concerns among DBAs and developers regarding their impact on the Db2 subsystem and applications.

Dispelling the myths and misconceptions surrounding these levels is our goal. Through this event, we aim to equip you with a deeper understanding of each level and how they synergistically function together.



**Leila Hosseini**  
Sr. DBA for Db2 on z/OS

**REGISTER**

### ENHANCING DB2 DATABASE DESIGN AND APPLICATION PERFORMANCE: A COMPREHENSIVE OVERVIEW OF DB2 12+ AND 13 FEATURES

Continued

Among the new features, one of the most exciting additions to Db2 13 is **SQL Data Insights (SQL DI)**, an AI-powered feature that combines deep learning and advanced IBM Z technologies to perform SQL-based semantic queries on user tables and views. This AI capability allows Db2 to identify data

relationships and patterns without moving the data outside of the engine. SQL DI comes with a Graphical User Interface and eliminates the need for ETL or complex machine learning models, making it a powerful and user-friendly tool.

Susan Lawson’s comprehensive overview of Db2 12+ and 13 features showcased the benefits and potential challenges of each enhancement. With the continuous evolution of Db2 and the introduction of new function levels, staying informed and conducting individual testing will be vital for optimizing database design and application performance. Leveraging these enhancements effectively will enable organizations to unlock the full potential of their Db2 databases, achieving enhanced performance, flexibility, and efficiency in their operations.

**YI&A SQL Data Insights**

- SQL Data Insights (SQL DI) - Optional 13 Feature
  - Artificial Intelligence (AI) functionality in the Db2 engine
    - Identifies valuable hidden insight within a relational database for better business decisions
    - Reduces the time necessary to take advantage of AI
    - No extract-transform-load (ETL) processes to move the data off mainframe
    - No need for machine-learning models to be designed by expert data scientist
  - SQL DI user interface can be added a Db2 table or view as an AI object
    - Enable AI query functionality
    - Run semantic AI queries directly against the Db2 table or view
      - Providing insight into the data and make business decisions
    - Three SQL DI Db2 built-in functions for semantic queries
      - Used in queries to infer hidden relationships between different entities in a table or view

- AI\_SIMILARITY
  - Find similar or dissimilar entities in a table or view
- AI\_SEMANTIC\_CLUSTER
  - Check what other entities exist in table or view that could belong to a cluster of up to 3 entities
- AI\_ANALOGY
  - Determines whether the relationship of a pair of entities applies to a second pair of entities

Db2 12+/13 for z/OS: Database Design and Application Performance Features and Usage 67



**Access the Recording**



## NEWS AND ANNOUNCEMENTS

- Sponsorship opportunity—We are looking for additional co-sponsors for the Virtual Db2 user group. The user group has been in existence since 2022 and is gaining respect among users of Db2. The user group gives its sponsors an opportunity to show that they are working with, and helping to build, the Db2 user community. Contact [virtualusergroups@gmail.com](mailto:virtualusergroups@gmail.com) for more information.
- After you register for our next event on September 26, save the date for our November 21 session by IntelliMagic.

## ABOUT THE VIRTUAL DB2 USER GROUP

The Virtual Db2 user group is an independently-operated vendor-neutral site run by and for the mainframe Db2 user community. This is a mainframe Db2 information website, not in any way related to, sponsored, or approved by IBM, which is the legitimate owner of the trademark, and any use of the mark in the URL or the body of the site is for information, education, and opinion expression purposes. The Virtual Db2 user group was established as a way for individuals using IBM's Db2 for z/OS database to exchange information, learn new techniques, and advance their skills with the product. i Anyone with an interest in Db2 for z/OS is welcome to join the Virtual Db2 user group and share in the knowledge exchange.

## SPONSORS

