

**Db2 for z/OS Performance Updates**  
**40 years anniversary version**

**Akiko Hoshikawa**

*IBM*

# Disclaimer

© IBM Corporation 2023. 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.


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.

## Perception of English "r" and 'l' by Japanese speakers

Native speakers of Japanese who have learned English as adults have difficulty perceiving the acoustic differences between English /r/ and /l/


Row Level Locking will be Low Level Locking by a native Japanese

[https://en.wikipedia.org/wiki/Perception\\_of\\_English\\_/r/\\_and\\_/l/\\_by\\_Japanese\\_speakers#:~:text=Goto%20\(1971\)%20reports%20that%20native,extended%20periods%2C%20and%20can%20articulate](https://en.wikipedia.org/wiki/Perception_of_English_/r/_and_/l/_by_Japanese_speakers#:~:text=Goto%20(1971)%20reports%20that%20native,extended%20periods%2C%20and%20can%20articulate)



# Db2 for z/OS

Celebration of History



June 7, 1983  
Announcement of  
Database 2 aka DB2

- 1983 Version 1
- 1988 Version 2
- 1993 Version 3
- 1995 Version 4
- 1997 Version 5
- 1998 Version 6
- 2001 Version 7
- 2004 Version 8
- 2007 Version 9
- 2010 Version 10
- 2013 Version 11
- 2016 Version 12
- 2022 Version 13

# Db2 Performance - Now and Then (14 years ago)

## IBM Relational Warehouse Workload - Deliver transaction through JDBC T4

- Total 160 SQL per transaction(30 update, 10 delete, 25 prepare/describe, 30 open, 20 fetch, 20 close)
- An application server is on the zLinux LPAR on the same CEC

---

	<b>Aug 2009</b>	<b>Aug 2023</b>
Processor	z10	z16
z/OS version	V1R10	V2R5
Db2 for z/OS version	DB2 9	Db2 13

# Db2 Performance - Now and Then (14 years ago)

IBM Relational Warehouse Workload - Deliver transaction through JDBC T4

- Total 160 SQL per transaction(30 update, 10 delete, 25 prepare/describe, 30 open, 20 fetch, 20 close)

	Aug 2009	Aug 2023
Processor	z10	z16
z/OS version	V1R10	V2R5
Db2 for z/OS version	DB2 9	Db2 13
Number of SQL per transaction	160	160
Application response time (ms)	48.0	10.678
Db2 elapsed time (ms)	13.476	6.7
Db2 CPU GP+zIIPs time (ms)	4.031	1.239
Transaction rate (per sec)	23	97
Lock req per trx	67	59
Getpage per trx	198	158
Log record size * #log record	93 bytes * 164	220 bytes * 176

# Now and Then - Observations..

- Hardware + Software combind
- Transaction rate - 4.2x more
- Transaction response time – 4.5x less
- Db2 elapsed time – 2x less
- Db2 CPU time - 3.2x less
- Getpage – 20% less
- Db2 logging – 2.5x more



# Agenda

- **Now and Then**
  - zIIPs
  - Data sharing
  - Open Data Set
- **Performance updates and Db2 13 expectation**
  - Db2 13 performance and expectations
  - Greater operational insights



# Db2 usage of IBM z Integrated Information Processors (zIIPs)



## DRDA access over TCP/IP

Distributed threads through running under Enclave SRBs (up to 60%)

including native SQLPL, native REST call over HTTP



## Parallel query

Child task of query parallelism with Degree > 1 (up to 100%)



## Utility processing

Portion of LOAD, REORG and REBUILD, 100% RUNSTATS



## XML processing

Schema validation and XML parsing (up to 100%)



## Db2 system agents

Db2 system tasks running under DBM1, MSTR address spaces except p-lock negotiation (up to 100%)

Index parallel update, pseudo deleted entry cleanup, FTB, etc.



## SQL Data Insights

Training process (up to 100%)  
AI built in functions (up to 100%) - internal use of parallelism processing



## Others

Db2zAI collection and learning  
Insync log read



# IBM Z Hardware Synergy

## Data sharing enhancements



### zEC12

- Large frames
- zEDC LOB compression
- Write around (CFCC 17)




### z13

- **Async CF duplexing (CFCC 21)**
- IDAA on zLinux
- SMT-2
- SIMD



### z14

- Encryption performance
- Huffman compression (entropy encoding)
- zHyperLink READ/WRITE
- CF encryption (CFCC 22)
- **Async Cross Invalidation (CFCC 23)**
- 25Gb RoCE Express2
- 32TB memory



### z15

- On chip compression to replace zEDC for large data (LOB, VSAM)
- Z SORT accelerator
- System Recovery Boost – with z/OS IPL
- Privacy for Diagnostics
- 40TB and TLB growth



# IBM z16 and Db2

- **Subsystem level recovery boost**
  - Middleware region startup boost in WLM boost attribute in STC classification
  - SVC dump boost
- **CFCC 25 (z16) short link improvement**
  - Faster CF operation => data sharing cost reduction
  - Note: strongly recommended to use CF sizer to reevaluate small structure size (<100MB) as we observed increase in size and/or reductions of entries.
- **GBP residency time reporting**
  - Db2 13 V13R1M100 above
  - IXLCACHE READ\_STGSTATS returns weighted average residency time in microseconds for directory & data
  - Note : This is applicable only if there are reclaim
- **SQL Data Insights**
  - Db2 requests AI operation acceleration through z/OS AI Optimization library (ZAIIO)
  - Note : z14, z15 & z16 above utilizes SIMD by OpenBLAS vector optimization
  - AI\_ANALOGY query utilize z16 AI on-chip accelerator

# GBP Residency Time (Db2 13)

DSNB783I -DBCL CUMULATIVE GROUP DETAIL STATISTICS SINCE 19:10:52 MAY 3, 2022

DSNB784I -DBCL GROUP DETAIL STATISTICS 750  
READS

:

DSNB787I -DBCL RECLAIMS 753

FOR DIRECTORY ENTRIES = 517522

FOR DATA ENTRIES = 1234861

CASTOUTS = 1499369

DSNB788I -DBCL CROSS INVALIDATIONS 754

DUE TO DIRECTORY RECLAIMS = 536277

DUE TO WRITES = 674783

EXPLICIT = 0

DSNB762I -DBCL DUPLEXING STATISTICS FOR GBP26-SEC 755

WRITES

CHANGED PAGES = 2113601

FAILED DUE TO LACK OF STORAGE = 42

CHANGED PAGES SNAPSHOT VALUE = 10402

DSNB820I -DBCL AVERAGE RESIDENCY TIME IN MICROSECONDS 756

FOR DIRECTORY ENTRIES = 19984895

FOR DATA ENTRIES = 5446695

DSNB790I -DBCL DISPLAY FOR GROUP BUFFER POOL GBP26 IS COMPLETE

Requires ..

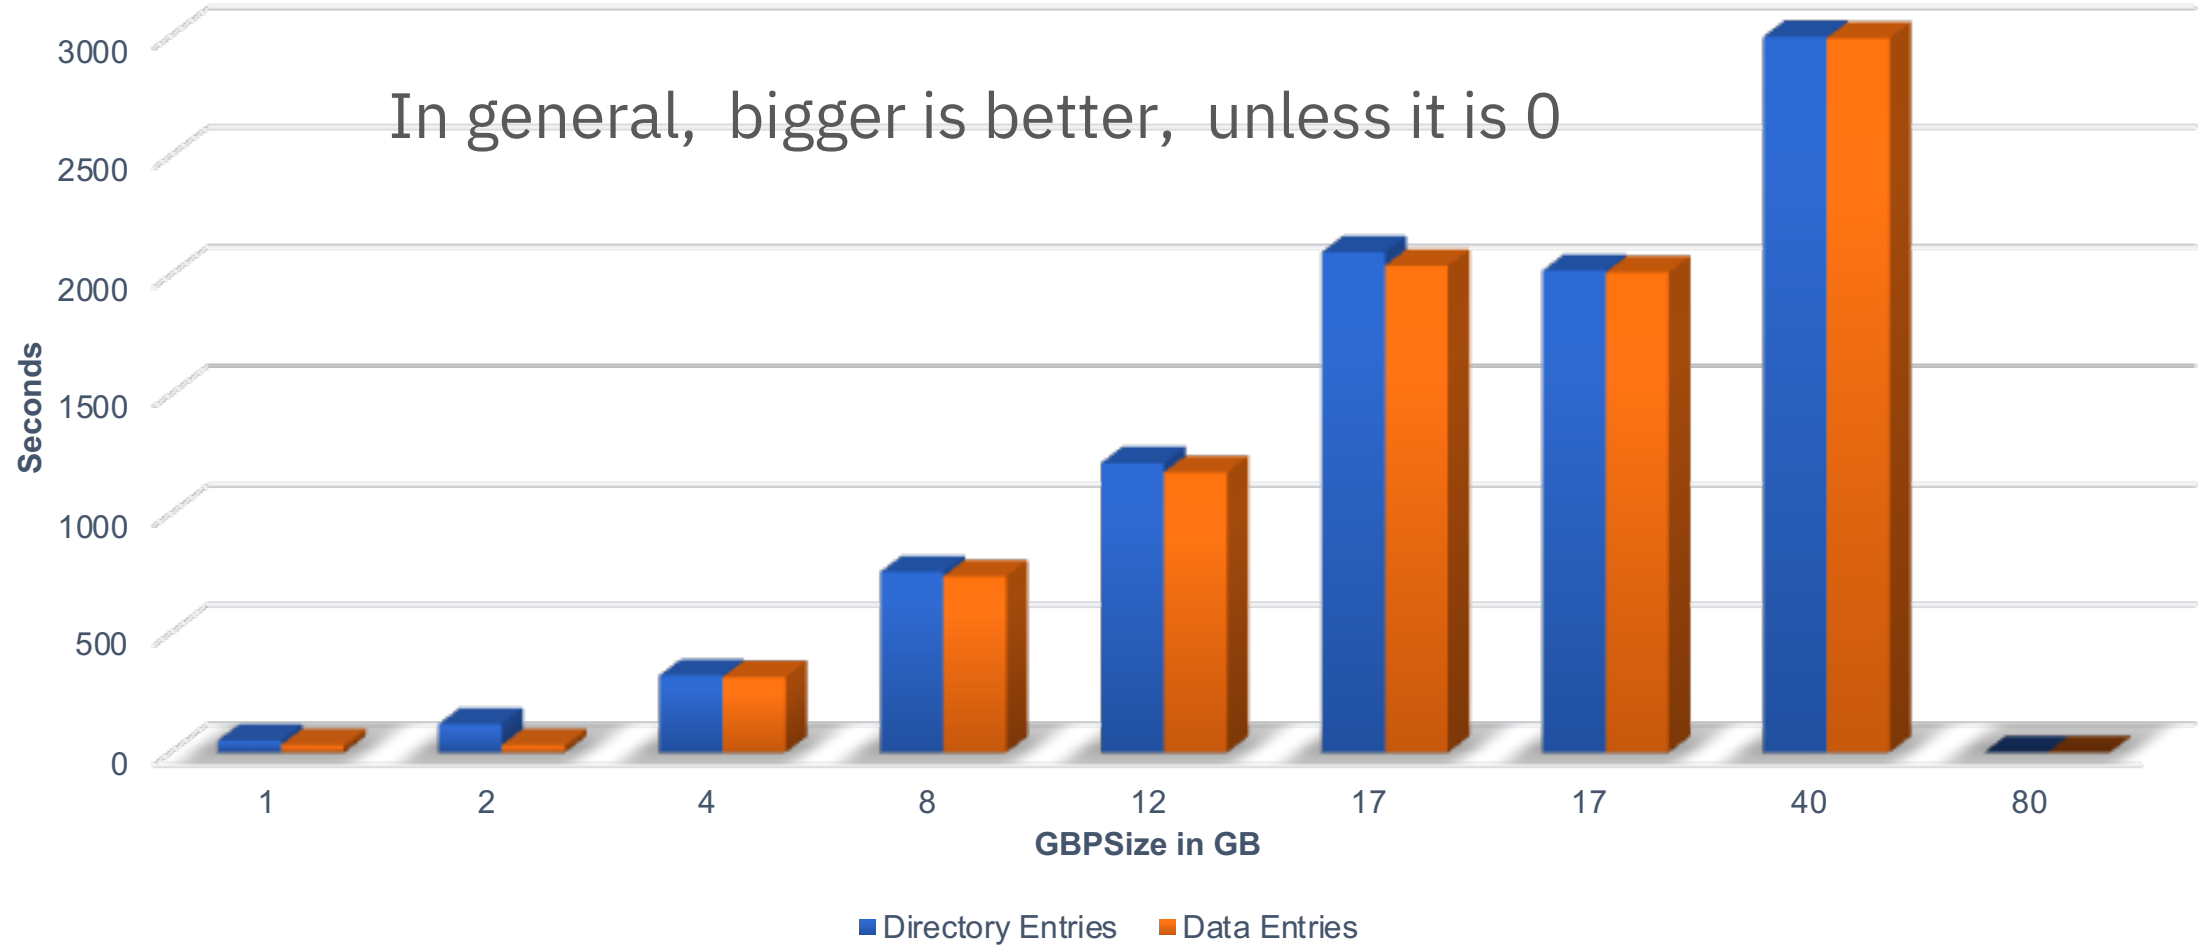
- CF on CFCC level 25 (z16)
- z/OS OA60650
- Db2 13 FL100
- Start measurement at reclaim.
- Without reclaim, residency time is 0
- No benefit of increasing GBP size once there is no reclaim

Published through,

- DISPLAY  
GROUPBUFFERPOOL  
command
- IFCID 230/254

# GBP Residency Time Example (Db2 13)

GBP Average Residency Time as GBP size increases - OLTP



# Max of DSMAX - History



# Open/Close Data Set Optimization

## – z/OS 2.5

- General reduction of below the bar (BTB) usage: 20-30% reduction seen of DBM1 BTB usage per data set
- Option to move Scheduler Work Blocks(SWBs) to above the bar (ATB) : additional 5-7% reduction per data set
  - PARMLIB (ALLOCxx) SYSTEM SWBSTORAGE(ATB)
  - Db2 13 checks the setting at start up and utilize during data set open

## – Db2 12 and 13 : close the data set only used by utility

- Pain-point : COPY utility requires large numbers of data set to be open and impact DSMAX for other important work
- Db2 12 APAR PH27493/PH33238 :
  - Monitor data set usage
  - During PCLOSET process, close data sets that were only accessed by utility processing
  - When DSMAX is hit, close data sets that were only used by utility first

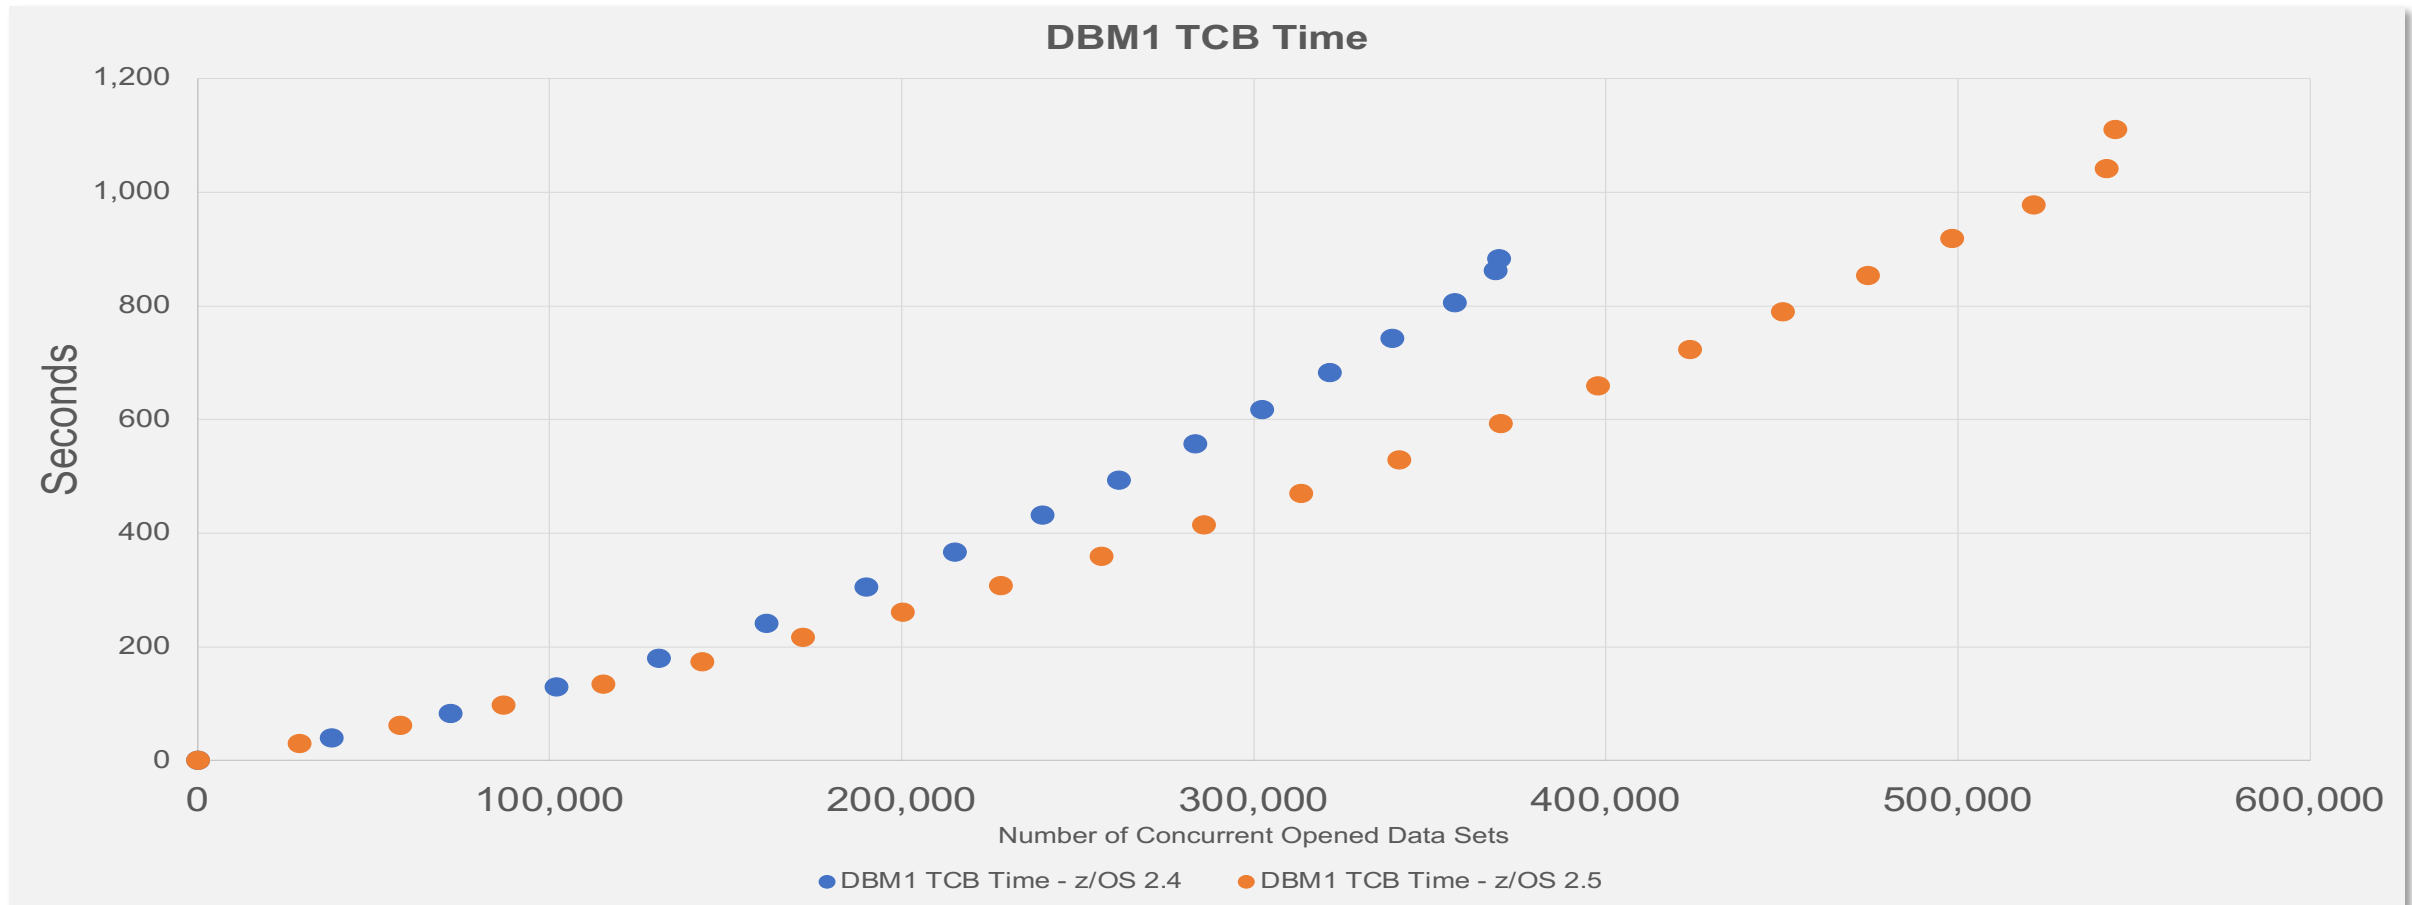
# Concurrent open data set : Theoretical DSMAX

z/OS 2.5 : significant reduction of DBM1 below the bar memory usage per data set

- Db2 can open more data sets in Db2 12 / 13 and faster than with z/OS 2.4

z/OS 2.5 + Db2 13: takes advantage of **SWBSTORAGE=ATB**

- Db2 13 can open even more! DSMAX zparm upper bound was updated to 400K



# Db2ZAI

Performance optimization based on what we learned + AI

```
BROWSE      DB2INT.AI16.BASE.SCOYSAMP(COYAKIRA)      Line 0000000000 Col 00
Command ==> _____ Scroll ==>
***** Top of Data *****
//*****
//*  NAME = COYAKIRA
//*
//*  DESCRIPTIVE NAME = Proc for starting the Db2ZAI Liberty server
//*
//*  Licensed Materials - Property of IBM
//*  5698-CGN
//*  Copyright IBM CORP. 2020, 2023
//*
//*****
//COYAKIRA  PROC
//E1  EXPORT SYMLIST=(ZAIHOME)
//*-----
//S1  SET ZAIHOME='<DB2ZAI_HOME>'
//*-----
//* Proc for starting the Db2 for AI for z/OS Liberty Server (Akira)
//*
```



# Db2 13 and Performance Updates

# Performance Redbook

- Published in Jan, 2023
- <https://www.redbooks.ibm.com/abstracts/sg248536.html>
- Written by Db2 for z/OS performance team
- Covers major performance items delivered by Db2 12 Continuous Delivery
- Covers major IBM zSystems synergy items since last redbook
- Large variety of writing styles
- Comprehensive performance measurements
  - Db2 11 300 pages
  - Db2 12 380 pages
  - Db2 13 414 pages

## IBM Db2 13 for z/OS Performance Topics

Neena Cherian  
Nguyen Dao  
Reynaldo Fajardo  
Akiko Hoshikawa  
Peng Huang  
Maggie Ou Jin  
Ping Liang  
Jie Ling  
Todd Munk  
Bart Steegmans


Jasmi Thekveli  
Lingyun Wang  
Chung Wu  
Chongchen Xu  
Huiya Zhang  
Xiao Wei Zhang  
Xue Lian Zhang



 Analytics

Data and AI

# IBM zSystems Synergy (Chapter 3)



## 100+ pages to cover IBM Z synergy items since last redbook

1. Db2 workloads performance comparison z15->z16
2. Data sharing cost reduction using z16 CF short link
3. CFLEVEL 25 migration consideration for small structures
4. Study of GBP residency time and increasing GBP size (Db2 13 only)
5. Study of system recovery boost (zIIP boost only)
6. Exploitation of Z Sort by REORG utility
7. Exploitation of Z Sort by Db2 SQL sort
8. LOB compression using z15 integrated accelerator for zEDC (on-chip compression)
9. Async Cross invalidation with and without distance
10. zHyperLink read and write support
11. Huffman compression
12. Study of data set and CF structure encryption



---

Highly available  
system management  
with greater resiliency  
and flexibility

Scalability and  
performance  
improvement through  
smarter optimization  
without needing tuning  
actions

New business  
insights using Deep  
Learning without  
complex AI  
application  
deployment

Produce critical  
operational  
information in  
managing complex  
enterprise systems

# Db2 13 Functions at GA (FL100, FL500, and FL501)

## 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
- 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

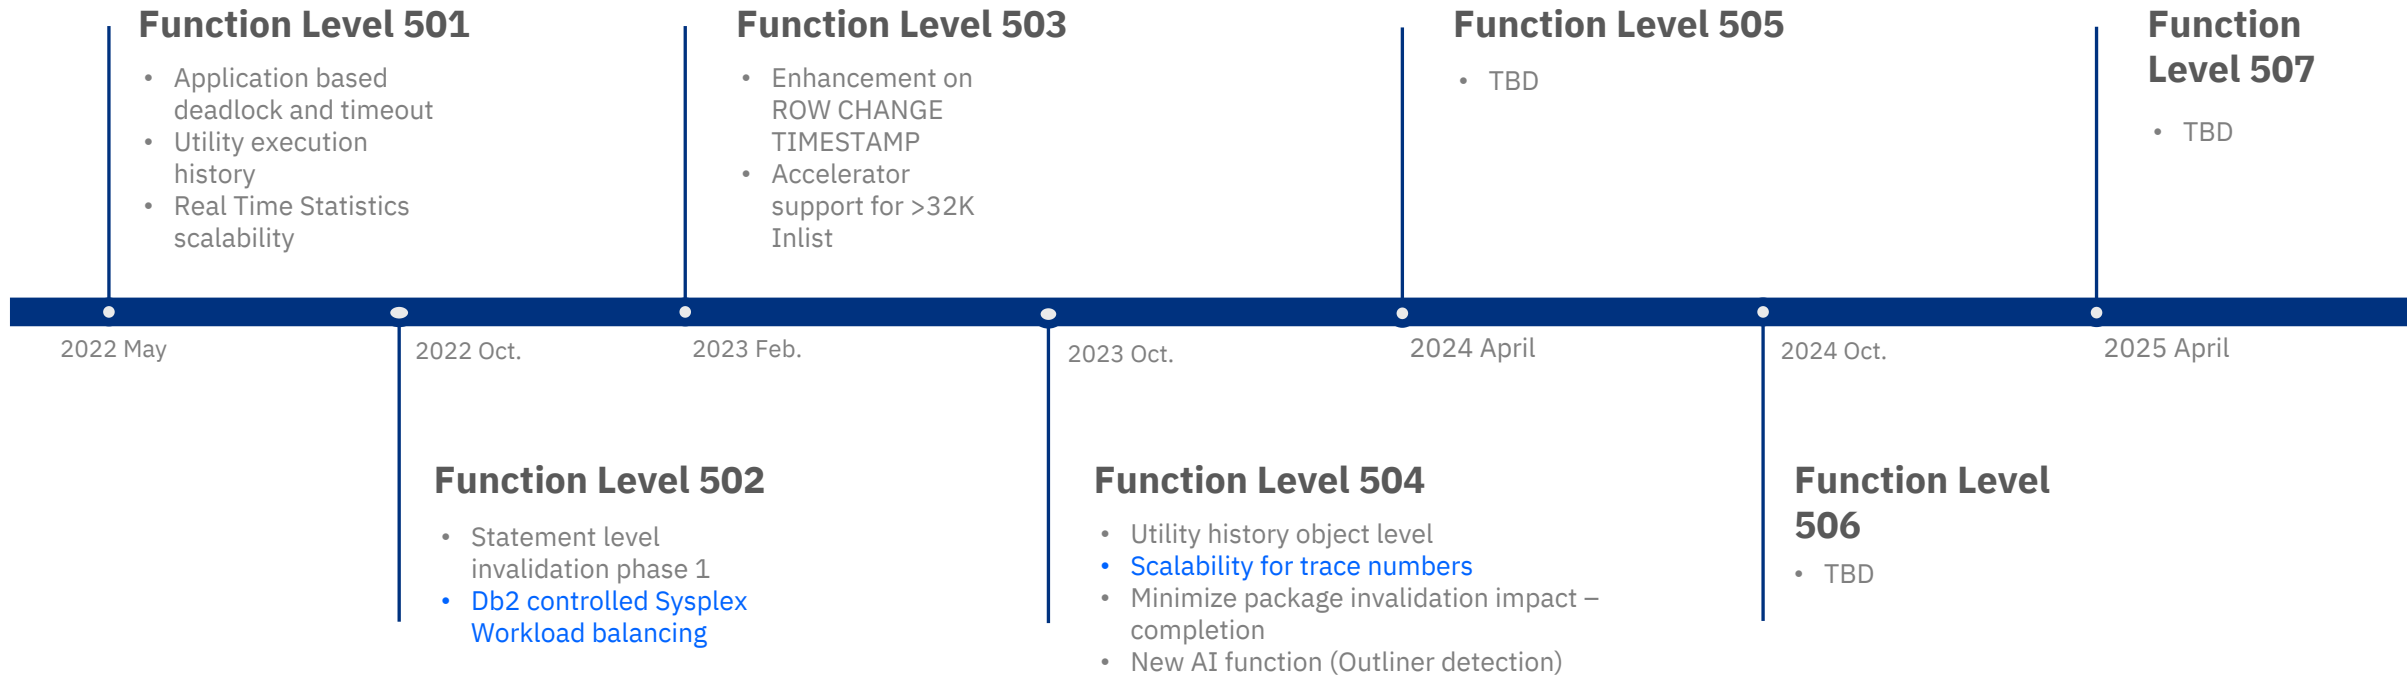
## V13R1M500 (FL500)

- Increased flexibility for package ownership
- Page sampling for inline statistics
- SQL Data Insights
- Reduced ECSA storage for IFI buffers Part2
- 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

# Building Upon Db2 13



## On top of GA Function Level - 2022

- Enhancement in profile monitor thread connection capability
- Db2ZAI V1.5 with granular filtering

## 2023


- Discovery and control for secure TLS connections
- Removing Stacking limitation for PBG to PBR conversion
- Performance enhancement using vector prefetch in SQL\_Data Insights
- PBR relative scalability retrofit in V12
- FTB monitoring enhancements
- Db2ZAI V1.6 with simplification

## 2024 (to be committed)


- SQL Data Insights incremental training
- Further schema flexibility
- Utility enhancements
- GBP invalidation enhancements
- And more

# Migration (CATMAINT, Activate) performance

Migration elapsed and CPU time (sec) with small system from Db2 10 to 13



# Db2 13 Performance Expectation



## Sequential Batch Updates

- Up 10% CPU reduction in sequential insert, update or delete operations compared to Db2 12 due to smarter index look aside



## OLTP workloads

- Equivalent performance if same zparm setting is used
- A few % CPU reduction if default FTB zparms are used due to increased FTB eligibility
- Noticeable CPU reduction for random insert across PBG partitions



## Query workloads

- Mostly equivalent performance
- Possible improvement with selective sort intensive queries



## Utility

- Equivalent except..
- Up to 60% CPU reduction REORG INDEX with default zparm.
- Inline stats to utilize page sampling in REORG / LOAD



# Db2 13 Improvement


## OLTP (or mainly random access )

- **More index to be eligible for FTB**
  - Range of improvement in random access : 0 to 5%
    - Default of FTB\_NON\_UNIQUE\_INDEX subsystem parameter is updated from NO to YES (V13R1M100)
    - Index key size increase (V13R1M500)
  - ✓ Benefit : Getpage decrease, class 2 CPU time reduction
  - ✓ Trade-off : DBM1 64 bit private memory increase, IRLM notify/CPU increase
- **Smarter cross partition search while inserting to PBG (M100)**
  - Range of a few% to 80% improvement in random concurrent insert with page level locking against PBG
- **Smarter 64 bit storage management (M100)**
  - RSM=AUTO (V12) users have reported noticeable MSTR non-zIIP SRB reduction in V13

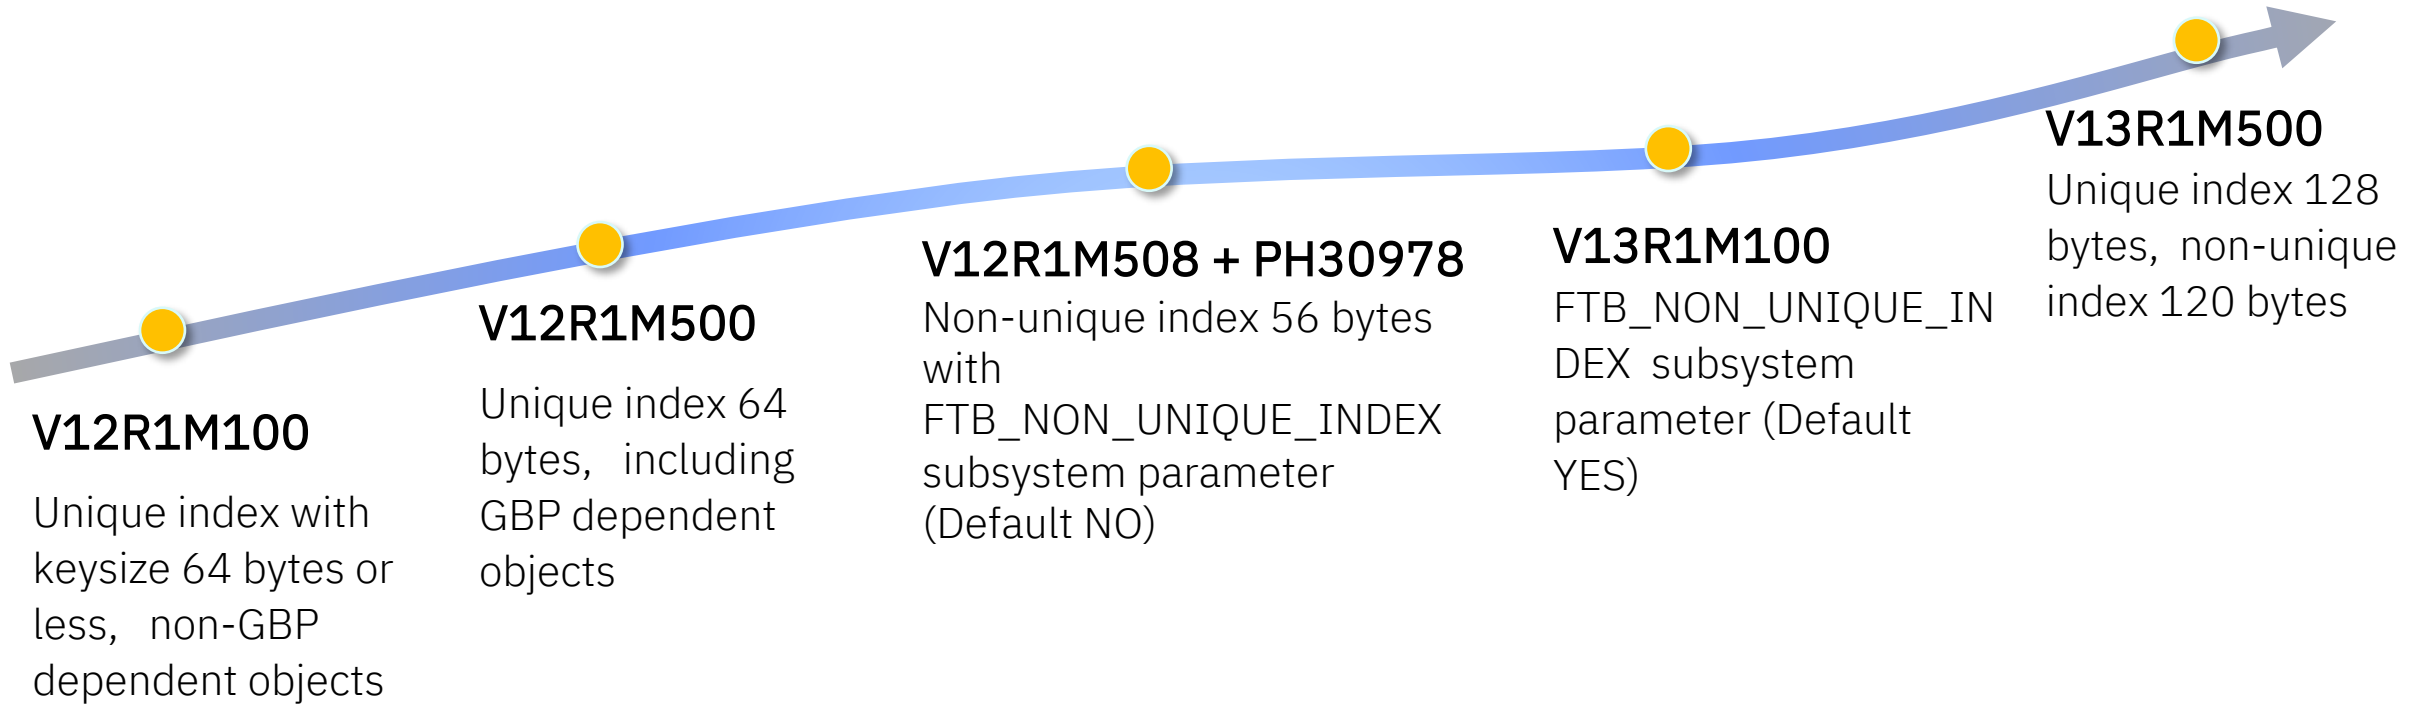
## BATCH (or mainly sequential access)

- **Index Look Aside for Insert/Update/Delete**
  - Detect the access pattern from the execution and enable index look aside for SQL INSERT/DELETE/UPDATE without relying on the RUNSTATS information
  - Work seamlessly with FTB enabled/disabled
  - Up to 10% improvement (range observed was +1 to -19%)

CPU improvement in (%) in Db2 13 on Sequential Update against Db2 12



# Journey of FTB





## FTB selections in INDEX\_MEMORY\_CONTROL

- **AUTO (Default)**
  - Selecting from higher index traverse count during monitoring interval
- **SELECTED, AUTO (or, n) option by PH23238**
  - Limit the candidates to the list provided in SYSIBM.SYSINDEXCONTROL with ACTION=A

IBM brokerage  
OLTP workload

Function Level




# Db2 13 Improvement Query

- **ORDER BY and GROUP BY sort (V13R1M100)**
  - SORT related zparm default updates
    - SRTPOOL default : 10MB to 20MB
    - MAXSORT\_IN\_MEMORY: 1M to 2M
  - More Z Sort
  - Combined with various internal optimization, up to 10% CPU reduction is observed from sort intensive queries

## Utility

- REORG\_INDEX\_NOSYSUT1 subsystem parameter default is updated from NO to YES (M100)
  - Up to 80% elapsed time and 50% CPU reduction
  - Minor memory usage increase seen (50MB)
- INLINE statistics for REORG / LOAD to utilize page level sampling by TABLESAMPE SYSTEM AUTO. (M500)
  - Up to 20-70 % CPU (mostly zIIPs) reduction  
(More details in Haakon's session)

# Db2 13 Performance Examples



# Storage Management

## Above The Bar (ATB)

### Db2 12 FL100 : REALSTORAGE\_MANAGEMENT = AUTO

- Db2 releases 64 bit real memory at thread level
- DBATs are terminated at every 200 commit
  - APAR PH36114 limits the max number of pooled disconnected DABT terminations to 50 per cycle
- Recommendation : If you have enough available real memory, use RSM - OFF
  - APAR PH47163 implements the recommendation avoid discard if  $\text{MIN}(\text{REALAVAIL}) > 1.2 * \text{MAXSPACE}$  with RSM=AUTO
  - New value for RSM=AUTO1 is added to support existing behavior

**Note : if your Db2 is running on IBM z16 with RSM=AUTO/AUTO1, install z/OS 2.5 OA61261, z/OS 2.4 OA64487 to reduce the impact from DiscardData operation**

# Storage Management

## Above The Bar (ATB)

### Db2 13 FL100 :

- Deprecated REALSTORAGE\_MANAGEMENT zparm.
- Instead of thread driven discard request, a system thread will monitor the storage usage and issue the discard from the largest free slots.
  - Note : DBATs are terminated at every 500 commits
- Further optimization in DBATs termination in POOLINAC processing
- In a well-behaved environment, released storage will be reused by the subsequent threads. Expect to have similar performance of RSM=OFF in Db2 12
- Db2 13 may see higher HWM of 64-bit memory if there are surge of thread creations
- RSM=AUTO (V12) users have reported noticeable MSTR non-zIIP SRB reduction in V13

# Greater Operational Insights

Information that helps to manage complex enterprise systems



# Utility History

- Db2 V13R1M501
- UTILITY\_HISTORY zparm
  - UTILITY | NONE
- New catalog table SYSIBM.SYSUTILITIES
  - A row per each utility execution
  - NAME of utilities
  - STARTTS and ENDTS
  - ELAPSED TIME, CPU TIME and ZIIPTIME
  - RETURN CODE, CONDITION
  - NUMOBJECTS
  - STARTLOGPOINT
- Validate utility executions / failures
- Historical trend to balance the workloads

## Utilities that are recorded

- BACKUP SYSTEM
- CATMAINT
- CHECK DATA
- CHECK INDEX
- CHECK LOB
- COPY
- COPYTOCOPY
- LOAD
- MERGECOPY
- MODIFY RECOVERY
- MODIFY STATISTICS
- QUIESCE
- REBUILD INDEX
- RECOVER
- REORG
- REPAIR
- REPORT RECOVERY
- REPORT TABLESPACESET
- RUNSTATS
- STOSPACE
- UNLOAD



# Serviceability Updates

## Speed-up analysis

### Index Split Diagnostic Enhancement

- Db2 13 makes diagnostic of problematic index split easier
- RTS (SYSIBM.SYSINDEXSPACESTATS) reports index splits
  - REORGTOTALSPLITS number of index split
  - REORGSPPLITIME accumulated elapsed time for index split
  - REORGEXCPLITS number of index split that took more than 1 sec
- IFCID 369 (statistics, exception base) reports splits > 1 sec
  - Pageset, partition, page being split and URID (connect with log)
  - Total elapsed time

### Long lock and latch waiter

Pain-point : When there is an excessively long lock/latch wait, diagnosis requires recreate with additional documentation (trace or slip-dump)

### IFCID 3 (Accounting, thread level) records

- Information of longest lock/latch wait, service task wait, page latch wait from the record
- Speed up the analysis, possibly without additional documentation



# Instrumentation Updates

Item	Db2 12	Db2 13
STATIME_MAIN subsystem parameter to control interval of subset of statistics (IFCID 1,2,202,225,230,254,369,411,412 )	Y - Default 60 second PH18658, 2020	Y - Default 10 second
IFCID 369 (aggregated accounting by conntype) is moved (copied) to statistics class 1	Y PH43916, 2022	Y
Granular DRDA statistics IFCID 411 ( DRDA client application) and IFCID 412 (DRDA client userid) statistics	Y PH40243, 2022	Y
z/OS Workload Interaction Correlator & Navigator (zWIC/zWIN) Db2 support, requires z/OS OA61924	Y PH39140, 2021	Y
DDF DBATs termination details are added (IFCID 1, statistics)	N	Y
Long index split exception IFCID 396 (Statistics)	N	Y
GBP Residency time IFCID 230/254 (Statistics)	N	Y with CFCC level 25 (z16)
Long waiter information IFCID 3 (Accounting)	N	Y
Latch class expansion (Statistics)	N	Y

# ZPARM Updates (Performance Impacting Subsystem Parameters)

Change	Updated zparms	Default (old)	Comments/Perf impact
New	UTILITY_HISTORY	NONE	FL501above to turn on: only one row inserted per each utility run, then updated during the execution. Very minor performance impact per utility execution
New	SPREG_LOCK_TIMEOUT_MAX	-1	Application can override with own lock time out value. No overhead observed using register. Using profile shows base profile overhead
MAX	DSMAX		Updated max from 200,000 to 400,000
Default	FTB_NON_UNIQUE_INDEX	YES (was NO)	Possible CPU saving, utiliat more memory (64 bit)
Default	PAGESET_PAGENUM	RELATIVE (was ABSOLUTE)	Recommend to REOG in Db2 13 until Db2 12 retrofit
Default	STATIME_MAIN	10 (was 60)	Better observability, possible increase of SMF volume
Default	OUTBUFF	100MB (was 4MB)	Better resiliency, log read/write performance
Default	SRTPOOL	20,000 (was 10,000)	Better performance for relatively large sort
Default	REORG_INDEX_NOSYSUT1	YES	Better performance. Only behavior for REORG INDEX with SHRLEVEL REFERENCE/CHANGE
Default	MAXCONQN	ON (was OFF)	Best practice
Default	MAXCONQW	ON (was OFF)	Best practice
Default	EDM_SKELETON_POOL	81920 (was 51200)	Better performance accessing plan/packages
Default	EDMDBDC	40960 (was 23400)	Better performance accessing DBD
Default	MAXSORT_IN_MEMORY	2000(was 1000)	Better performance using in-memory workfile
Default	NUMLKTS	5000(was 2000)	Best practice, reduce lock escalation
Default	NUMLKUS	20000(was 10000)	Best practice, reduce application impact
Default	STATPGSAMP	Page level sampling	New default for inline statistics for LOAD, REORG TS
Obsolete	REALSTRAGE_MANGEMENT		Use new management
Obsolete	AUTHCACH		Use 4096
Obsolete	PARA_EFF		Use 50

# 64 + 1 Latch Class

LATCH CNT	/SECOND	/SECOND	/SECOND	/SECOND
-----	-----	-----	-----	-----
LC01-LC04	0.00	0.01	1.80	0.00
LC05-LC08	0.00	1.37	0.00	0.00
LC09-LC12	0.00	0.52	0.00	0.14
LC13-LC16	0.00	43.28	0.00	0.00
LC17-LC20	0.00	0.00	0.48	0.00
LC21-LC24	0.00	0.00	13.05	182.32
LC25-LC28	1.64	0.04	0.21	0.05
LC29-LC32	1.59	1.83	0.01	0.00
LC33-LC36	0.00	0.00	0.00	0.00
LC37-LC40	0.00	0.00	0.00	0.00
LC41-LC44	0.00	0.00	0.00	0.00
LC45-LC48	0.00	0.00	0.00	0.00
LC49-LC52	0.00	0.00	0.00	0.00
LC53-LC56	0.00	0.00	0.00	0.00
LC57-LC60	0.00	0.00	0.00	0.00
LC61-LC64	0.00	0.02	2.18	0.00
LC254	3.28			

- Db2 13 support more latch classes
- Storage manager related latch class 31/32 have been moved to new latch class 62/63
- Depending on the tool support, it may look that these latches are gone, but that is not the case

# Before Migrating to Db2 13

## Memory usage increase

- Majority of memory increase is expected from possible ZPARM updates
  - FTB, EDM pools, OUTBUFF, SORT related and REORG INDEX NOSYSUT1
- Thread memory
  - Expect similar with Db2 12 except,
    - Possible HWM increase due to SM logic updates and DBATs reuse limit (200->500)
    - Possible memory usage decrease REBIND with APREUSE
- ECSA – with large numbers of trace record turned on, possible to see small reduction
  - New ECSA estimate is documented in technical overview redbook

## Notes on instrumentation

- Db2 statistics default interval (STATIME\_MAIN) has been updated to 10 seconds, consider to use the same interval to compare the performance

# Latest Insert Performance APAR and Insert Algorithm 2 (IAG2)

- HIPER PH56059 (V12/V13) : applies to insert algorithm 1
  - Addresses the occasional performance degradation with high getpages in unbalanced insert workload across data sharing members. TS type PBG (MEMBER CLUSTER YES, APPEND NO)
- Status : Insert Algorithm 2

Thank you!



Akiko Hoshikawa

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



