IDUG

2026

Sydney | March 16 - 18

AU Db2 TECH CONFERENCE

Db2 for z/OS: Trends and Directions

Sueli Almeida, *IBM*

Session Code: ####

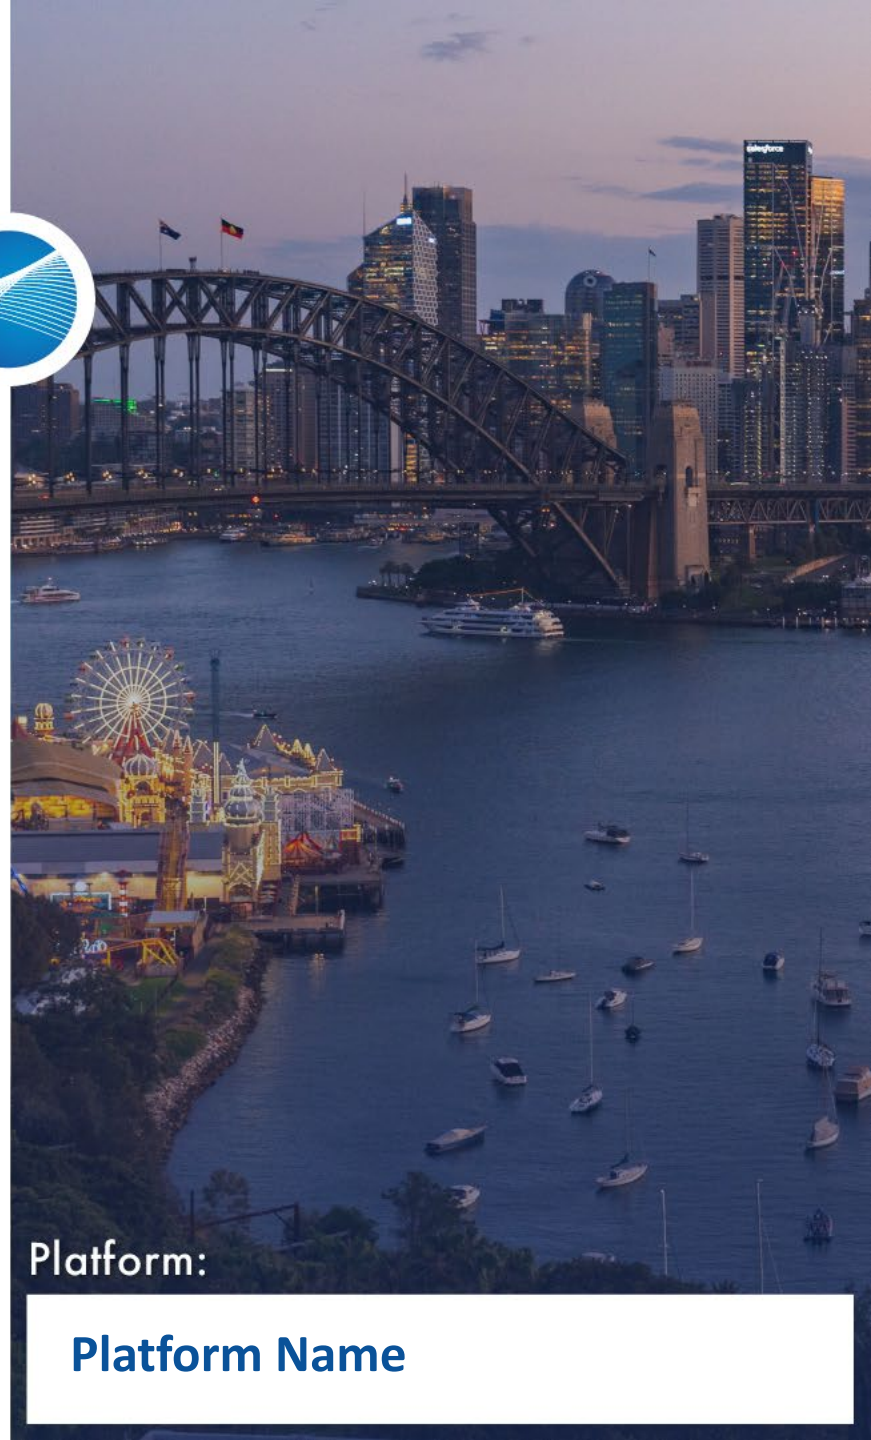

Platform:

Platform Name

Notices and Disclaimers

- 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.
- Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM.
- 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, the IBM logo, and ibm.com are trademarks of International Business Machines 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 the Web at “Copyright and trademark information” at: www.ibm.com/legal/copytrade.shtml.

Certain comments made in this presentation may be characterized as forward looking under the Private Securities Litigation Reform Act of 1995.
Forward-looking statements are based on the company’s current assumptions regarding future business and financial performance. Those statements by their nature address matters that are uncertain to different degrees and involve a number of factors that could cause actual results to differ materially. Additional information concerning these factors is contained in the Company’s filings with the SEC.
Copies are available from the SEC, from the IBM website, or from IBM Investor Relations.
Any forward-looking statement made during this presentation speaks only as of the date on which it is made. The company assumes no obligation to update or revise any forward-looking statements except as required by law; these charts and the associated remarks and comments are integrally related and are intended to be presented and understood together.

AGENDA

- Db2 for z/OS Technical Strategy
- Farewell Db2 12 for z/OS
- Db2 for z/OS Past, Present and Future
- Continuous Delivery & Function Levels update
- 2026 and beyond
- Db2 for z/OS – Exploit AI where it makes sense
- Summary

IBM Db2 for z/OS - Technical Strategy

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

Extend application and multi-tenancy capability.

Provide secure access to data wherever and whenever it is needed.

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.

SQL Data Insights

Db2 AI for z/OS & Db2 Agent

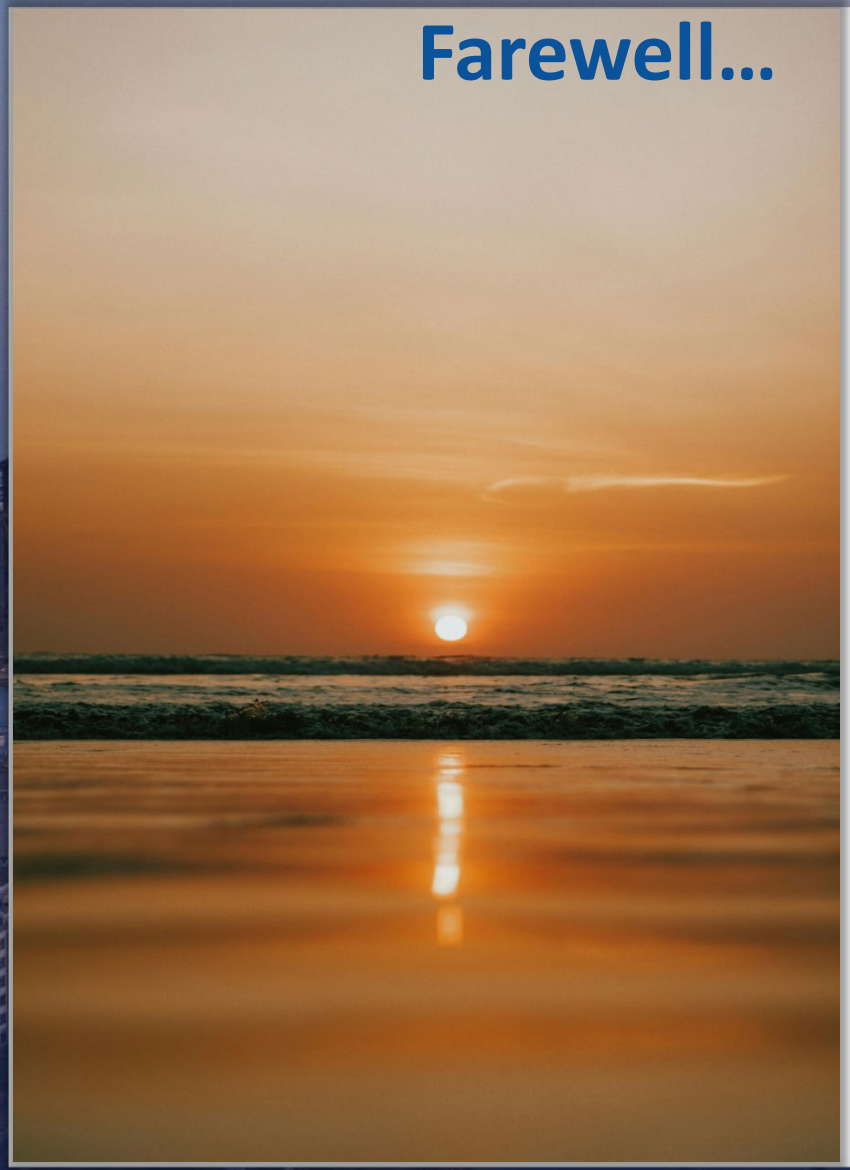

Farewell… Db2 12 for z/OS

- Do not forget about the PBR RPN red alert before migrating to V13!
  - Before migrating to V13R1M100 and before activating V13R1M500, read: https://www.ibm.com/support/pages/node/7158065
    - PH67878 to enforce red alert V13R1M100
    - PH66678 for pre-migration checklist
- Replace Data Studio with Admin Foundation and Developer Extension

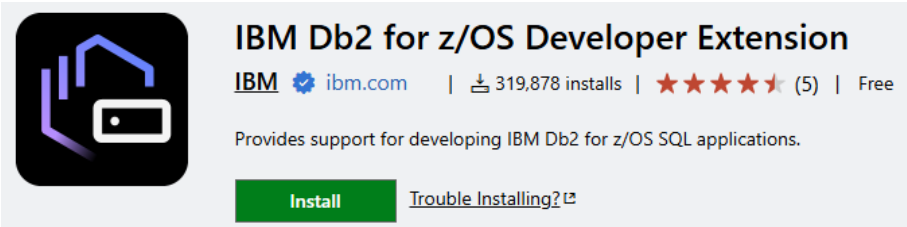

Db2 for z/OS
Past, Present and Future

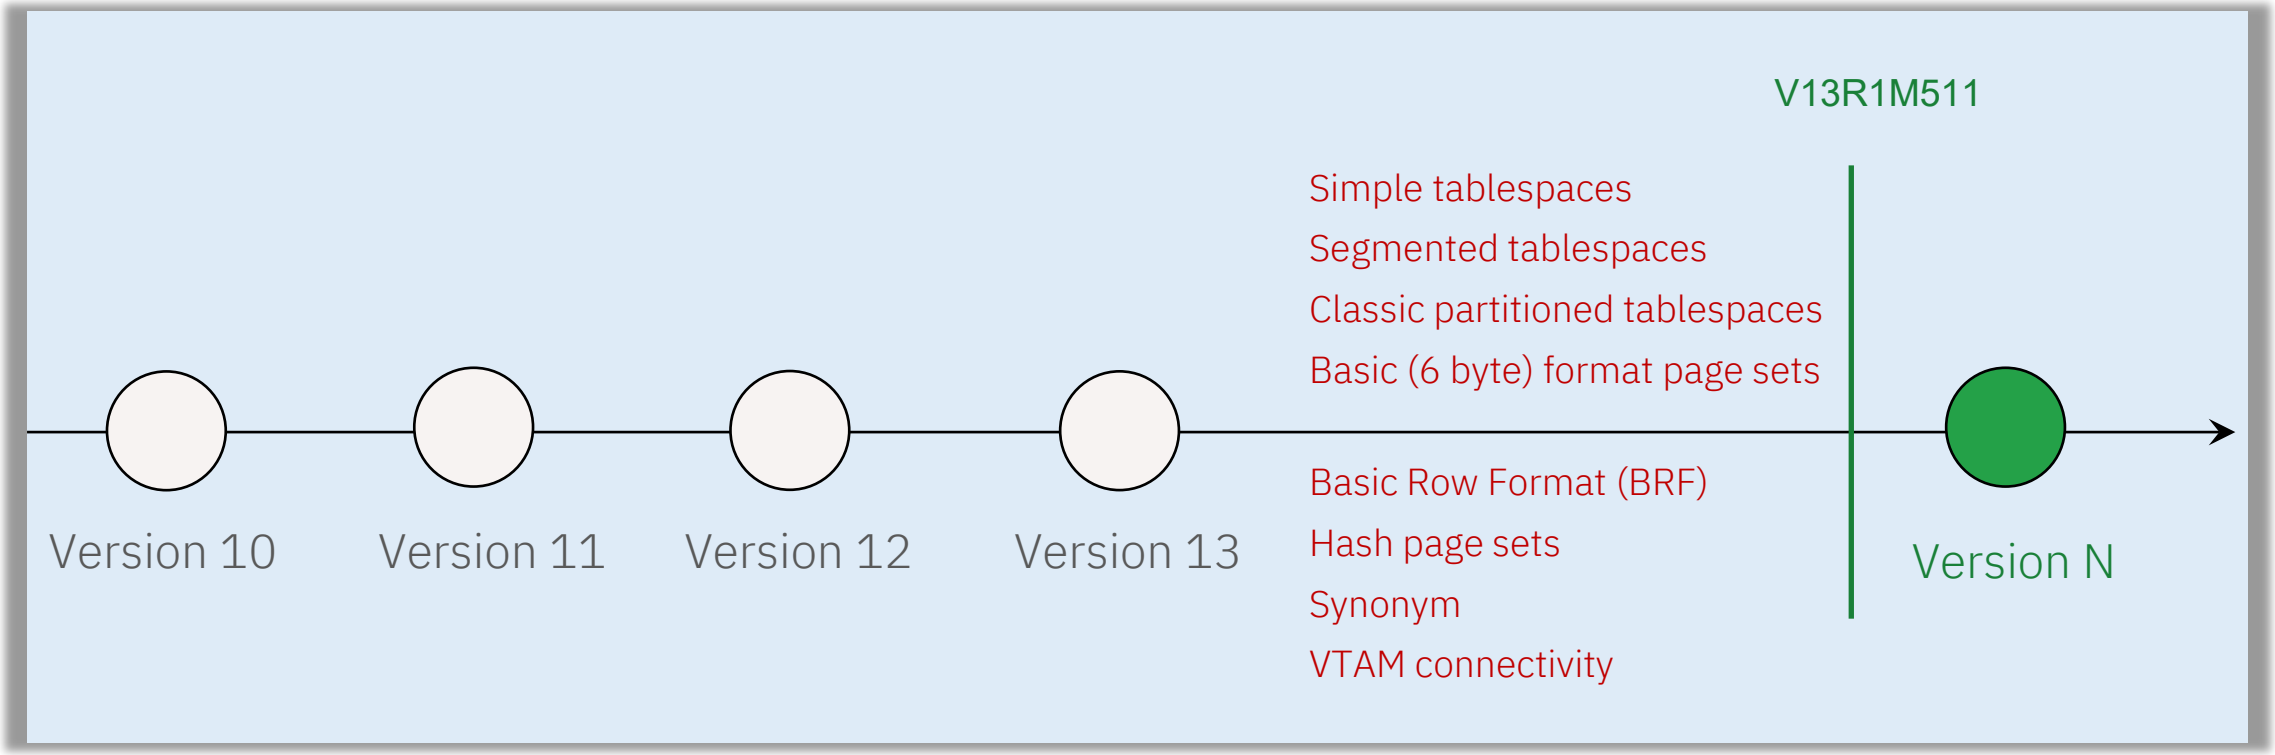

Blogs: Important news regarding Db2 for z/OS versions
Are you ready to convert from VTAM to TCP/IP in Db2 for z/OS?

Major Theme of Db2 Vnext

- New foundation for future innovation
- Simplification and automation
- Scalability and performance

Db2 for z/OS Release Plan

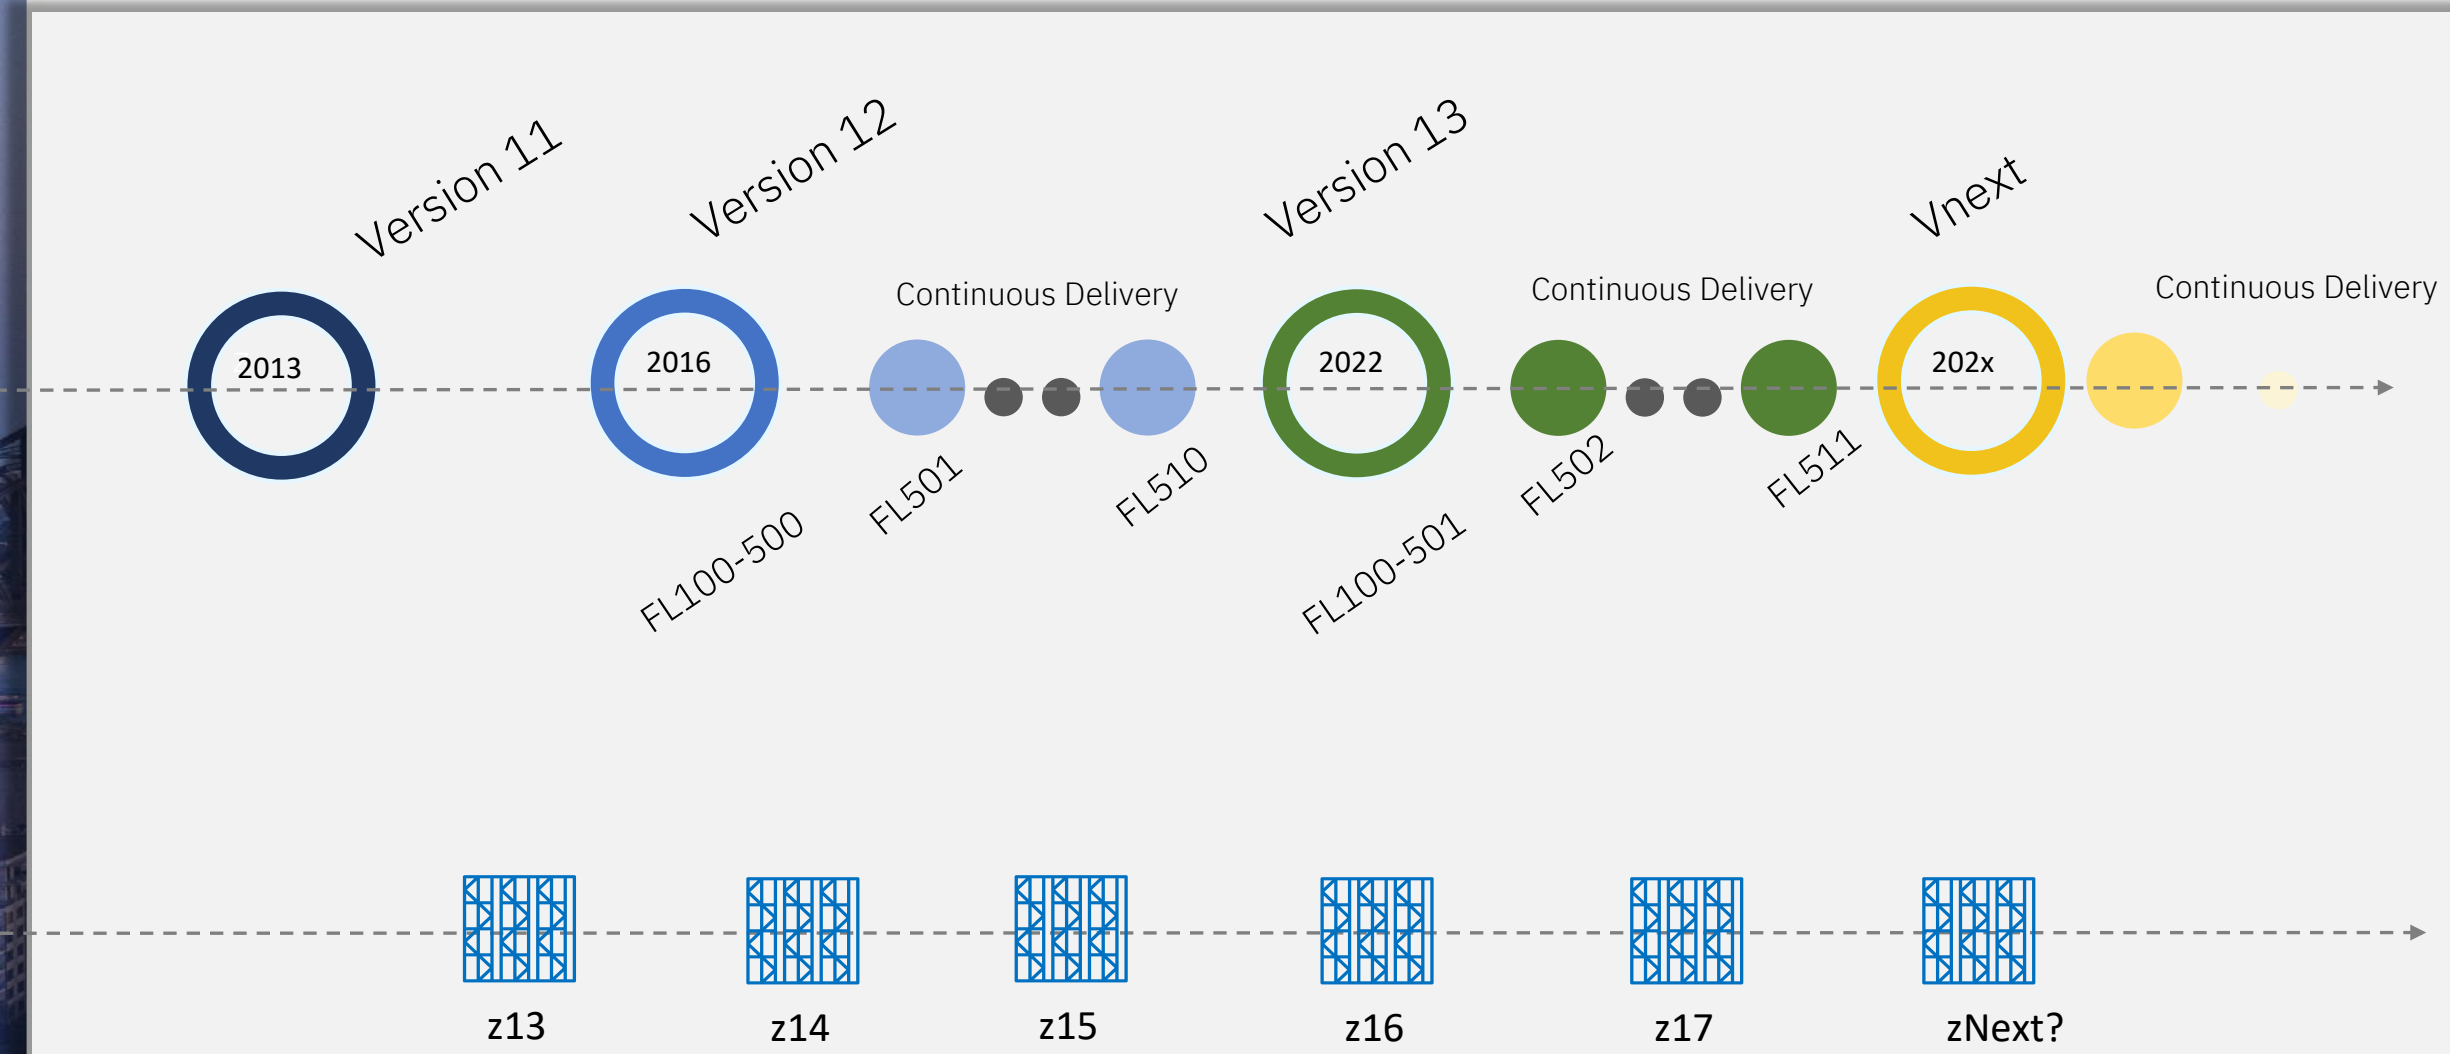

Db2 13 Continuous Innovation – GA (2022)

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

FL100...

- 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

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 MAXPARITIONS 254 by default

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 Continuous Innovation – Function Levels (1|2)

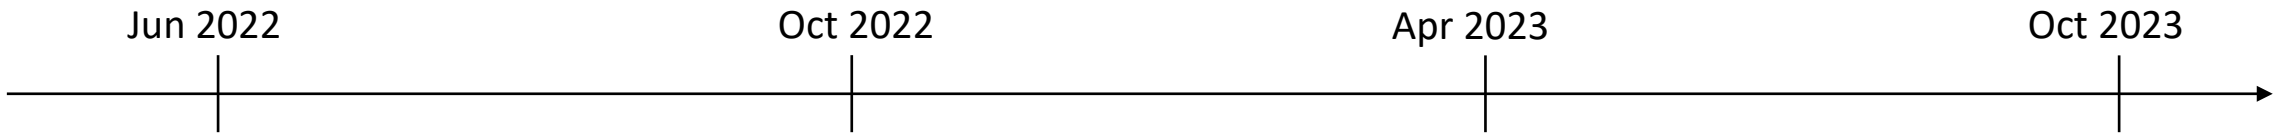

FL100, FL500, FL501

- Industry-first AI-infused insights into business data
- Enhanced operational flexibility with no application outages
- Improved interoperability with cloud applications
- Streamlined version upgrade with no application outages

And more…

FL502

- Statement-level dependency phase 1
- Db2-controlled sysplex workload balancing

FL503

- Default when adding row change timestamp columns
- Accelerator-only support for >32K in-list
- OPTIMIZE FOR n ROWS for SELECT INTO

FL504

- Object-level utility history
- Minimize package invalidation impact
- SQL Data Insights outlier detection
- More trace OP buffer destinations
- ORDER BY fullselect with LISTAGG

Db2 13 Continuous Innovation – Function Levels (2|2)

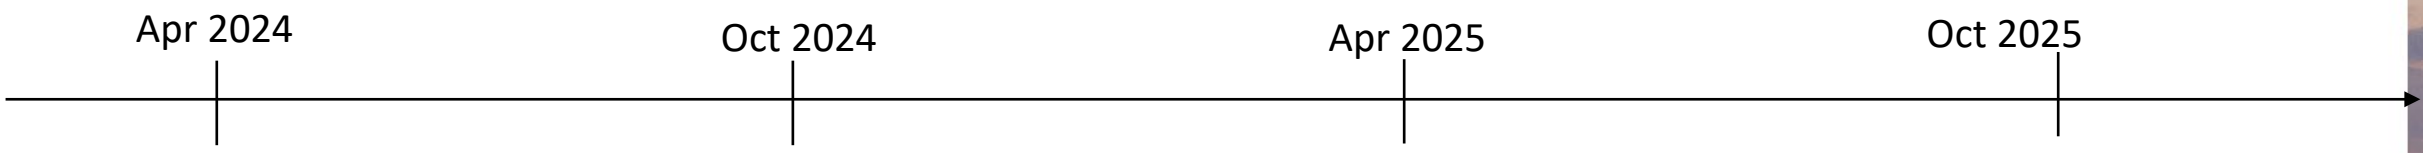

FL505

- RACF Identity Token(IDT) support
- Scalability for trace numbers
- Catalog security improvement
- INTERPRET scalar function
- Trace destination ZAI

FL506

- INLIST performance
- Column masking improvement
- >32 traces
- Multi-row insert with VALUES
- Alter add implicitly hidden ROWID column
- LOAD part REPLACE performance/availability through improved NPSI processing

FL507

- Online conversion from PBR to PBG & repartitioning
- System temporal concurrency
- LASTUSED col support in SYSPLAN
- Temporal and archive tables in a single SELECT
- >64Gb PRIQTY, SECQTY support

FL508

- Display Blocker enhancements
- Workfile management updates
- Generation of UUIDs
- REBIND phase in for SQL PL and trigger
- #of active log pairs increase to 1000
- Enhance local profiles

Db2 13 Continuous Innovation – Outside Function Levels (1|3)

2022

- New trace classes for distributed
- Improved –START/STOP ML command
- New LOAD option to reclaim NPSI space
- Improved messaging for utility FORCE option
- Improved –DISPLAY THREAD DBAT information
- Improved DECFLOAT support
- UNI_90 support in LOWER, UPPER, TRANSLATE
- List prefetch for MERGE
- PLAN_TABLE improvement for views & table expressions
- LOAD utility tape dataset estimation
- Accelerator support for LISTAGG, RAND
- PRDID values for specific DRDA levels
- Lock avoidance for singleton selects
- Control max number of concurrent UDFs to limit ECSA
- Remove stacking limitation for PBG to PBR conversions

2023

- Improved IFCID 0402 reset behavior
- Better index traverse counts in –DISPLAY STATS
- SQL Data Insights vector prefetch performance improvement & semantic clustering scoring accuracy
- APPLCOMPAT, PLANMGMT support for Db2-supplied routines
- Discovery and granular control for TLS connectivity
- Batch support for installation/migration CLIST
- SQLDA enhancement for REXX applications
- -DISPLAY GROUP migration readiness reporting
- Free inactive packages for native REST services
- Improve latch contention by reducing time to boost, and improved latch holder information
- Dynamic SQL Data Insights vector prefetch performance
- Simplified DROP VIEW authorization
- More granular accounting trace information for distributed threads
- Improved authorization for Db2 MQ tables
- BIND QUERY restriction removed for SQL PL packages
- Privilege improvements for –DIS PROFILE, -DIS DYNQUERYCAPTURE
- More responsiveness to FTB SYSINDEXCONTROL changes
- Separate location statistics from STATIME_MAIN
- Improved message handling for DSNJ110E log full conditions

Db2 13 Continuous Innovation – Outside Function Levels (2|3)

2024

- Improved message information for stale GBP pages
- Better profile options for monitoring connections for security
- Redirected recovery support for XML versioning
- CORS support for Db2 REST services
- Password phrase support in trusted contexts with JDBC Type 2 drivers
- SQL Data Insights model retraining
- Removal of ephemeral ports for DVIPA monitoring
- Improved DSNJ114I messaging for log archive errors
- DISPLAY privilege support for –DIS ML and –DIS RLIMIT
- New statistics for IP addresses accessing Db2
- SPECFICIMP support for WLM sysplex routing
- Bind location alias to different IP address for single subsystems
- ECSA relief by moving DDF WLM performance blocks
- Part-level support for MODIFY STATISTICS
- REBIND PACKAGE EXPLAIN(ONLY) support for advanced triggers
- COPY utility zIIP offload

2025 (January ~ June)

- Prevent console flooding by DSNL030I messages
- AUTHID based monitoring with security profiles
- EXPLAIN(ONLY) support for native REST services
- REST API support for SQL Data Insights administration
- Abbreviated APPLCOMPAT SQL syntax
- Improved sort workfile monitoring
- Datagram message support in MQListener
- Improved statement-level zIIP statistics
- IFCID350 dynamic SQL current schema name support
- DSNTXAZP improved discovery and update of DSNTIDxx members
- IPv4 and IPv6 subnet address support for trusted context ADDRESS
- Disable generation of GRANT statements for installation jobs
- Store decompression dictionaries in CDDS dataset to reduce IFI306, log read and GBP dependency overhead
- Zparm option to prevent MODIFY RECOVERY utility deletion of image copies
- Support for the same trusted contexts to be used for both local connections
- Support for a *queue* parameter in the CAF OPEN function
- New diagnostic level in the profile to reduce console messages
- LOB support for IDAA V8– store up to 2GB, result set up to 32K
- SQL Data Insights to be able to utilize secondary auth ID

Db2 13 Continuous Innovation – Outside Function Levels (3|3)

2025 (July~)

- Updates on ML commands to support imported data for Db2ZAI system assessment and SQL OPT dynamic filtering
- LISTDEF and Db2 installation CLIST enhancements for future catalog and directory conversions
- RACF group support in the trusted context SYSTEM AUTHID clause
- IDAA support for HASH built-in functions
- Enhancements in DSN1COMP for REORG option and for LOB compression
- Additional DSNV508I msg to warn 31-bit memory consumption in DIST address pace
- Improved control over new-connection requests for distributed applications
- Reduced z/OS console flooding by DSNL030I messages for long usernames
- Accounting information on whether a distributed client connection is secure
- Distributed tracing and OpenTelemetry support

https://www.ibm.com/docs/en/db2-for-zos/13.0.0?topic=13-2025-new-function-apars

2026 and beyond (1|2)

Core Db2 capability

- More granular application control using system profiles
- Scalability improvements
- More online schema => Online CREATE UNIQUE INDEX, DROP INDEX involved in RI, ALTER TABLE/ADD ALTER COLUMN, Transfer ownership of application objects
- Availability improvements => lower granularity for DBD lock; improve serialization on some internal operations)
- Db2 control over JDBC WLM parameters
- Db2 catalog improvements
- Cyber Vault enhancements
- Auto config discovery & Ansible playbooks, config as code for provisioning and migration
- Java 25 support
- and more …

2026 and beyond (2|2)

More Insights & Simplification

- SQL Data Insights incremental training & unstructured data
- Db2 AI for z/OS integration with SQL Tuning services & more user control
- Db2 AI for z/OS object level assessments
- DBA agent and watsonx Assistant for Z integration
- Utility history improvements
- RTS improvements
- Index traverse count stats
- Db2 command catalog table
- … and more

IBM Synergy, Db2 for z/OS and AI

Exploit AI where it makes sense

IBM Z Hardware Synergy

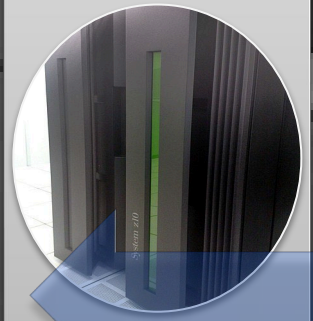

z10

- Large frames (1MB)
- PREFETCH H/W instruction
- HiperDispatch

z196

- Larger cache
- zHPF
- **Write Around** (CFCC 17)
- 3TB memory

zEC12

- Large frames (2GB)
- Flash Express
- zEDC LOB compression
- RoCE-express

z13

- **Async CF duplexing** (CFCC 21)
- IDAA on zLinux
- SMT-2
- SIMD
- 10TB memory

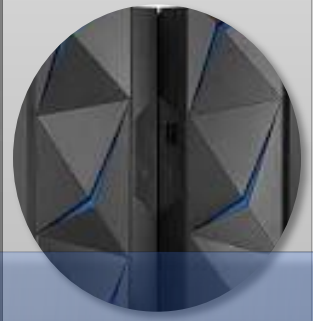

z14

- **Async Cross Invalidation (CFCC 23)**
- Encryption performance
- Huffman compression
- zHyperLink READ/WRITE
- CF encryption (CFCC 22)
- 32TB memory

z15

- On chip compression to replace zEDC for large data (LOB, VSAM)
- Z SORT
- System Recovery Boost
- Privacy for Diagnostics
- 40TB

z16

- **GBP residency time**
- **On chip AI accelerator**
- Subsystem level recovery boost
- CFCC short link improvement

z17

- remote zAIU accelerator
- 40% More cache
- more memory up to 64TB
- PCIe G4 for zHyperlink 2.0

IBM z17 and Db2 for z/OS

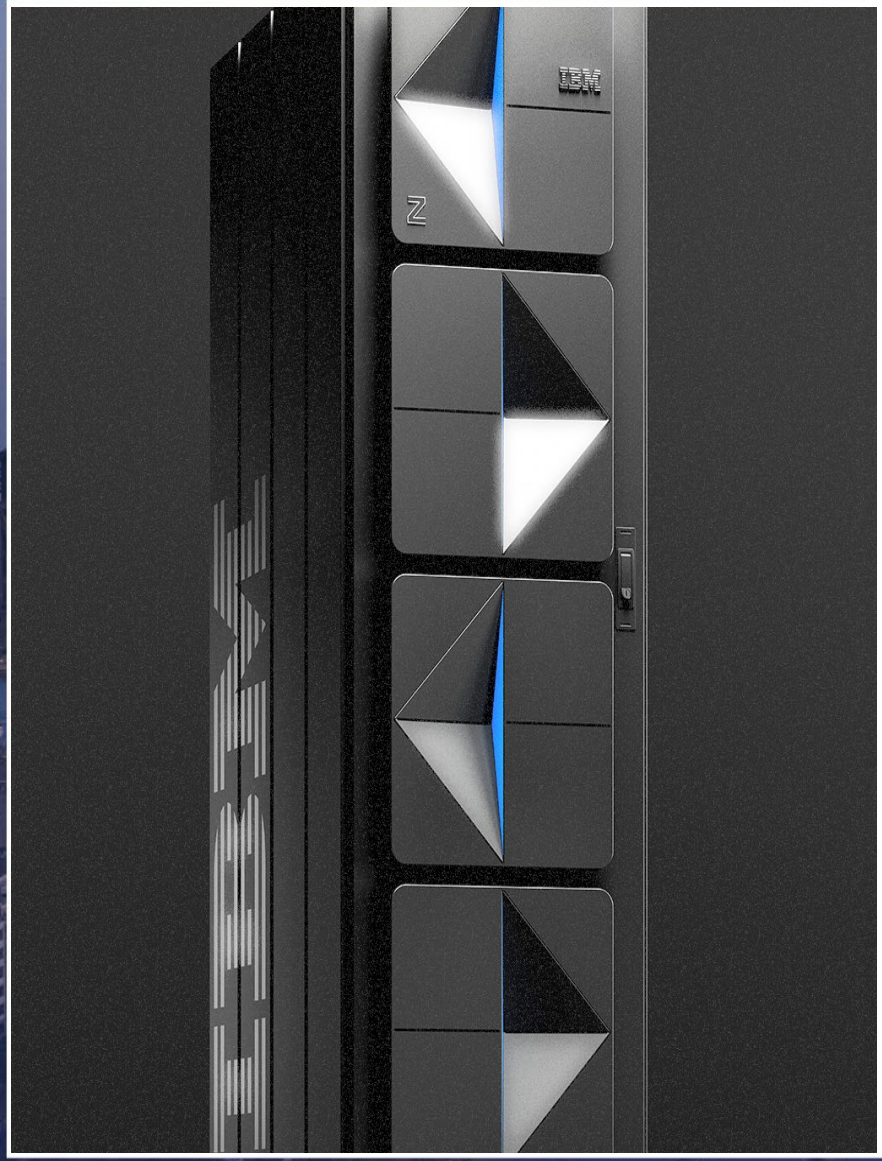

- **Base benefit existing Db2 workloads with z17**
  - 40% more cache per core over z16 to speed up Db2 workload
  - 64 TB max memory capacity from 48 TB
    - In-memory buffer pools PGSTEAL(NONE)
    - Thread reuse, REL(DEALLOCATE) and High performance DBATs
    - Fast Traverse Blocks (FTBs)
    - Insert Algorithm 2, so on
- **SQL Data Insights query acceleration with Telum II – 2nd gen AI on-chip accelerator**
  - Take advantage of intelligent routing of zAIU within the drawer
- **Combined with IBM DS8K G10, zHyperLink Express card which enabled PCIe Gen 4 improves Db2 synchronous read and write I/O significantly**
- **IDAA enjoys capacity increase**
  - **15%** over z16 on total capacity, **4%** more IFL per CPC
  - NVMe is part of Call Home
  - New Network Express Cards can be used for both OSA and RoCE requirements
- IBM Sensitive Data Tagging for Z

Blog: IBM z17: What It Means for Db2 for z/OS

IBM z17 and Db2 13 for z/OS

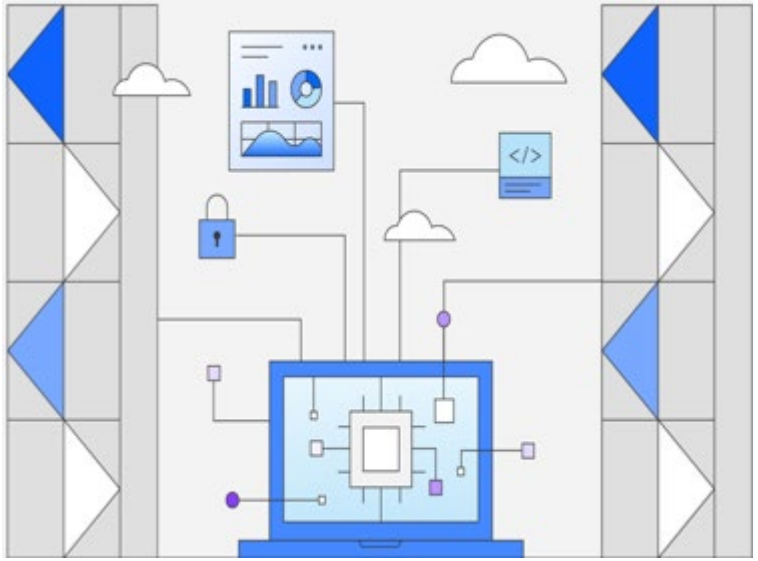

Statement of Directions & Delivery : related with Db2 associated with z17 and z/OS

- AI agents for middlewares (Db2, CICS, IMS, Concert, and IntelliMagic Vision) in support of IBM watsonx Assistant for Z
- Open Telemetry, distributed traces in middlewares (Db2, CICS, IMS, and MQ) improve observability across z middleware
- Sensitive Data Tagging on Db2 tables to discover and classify sensitive data
- Support of unstructured (text) data to perform AI_functions

Db2 SQL Data Insights

Unveil New Business Insights Through SQL Data Insights

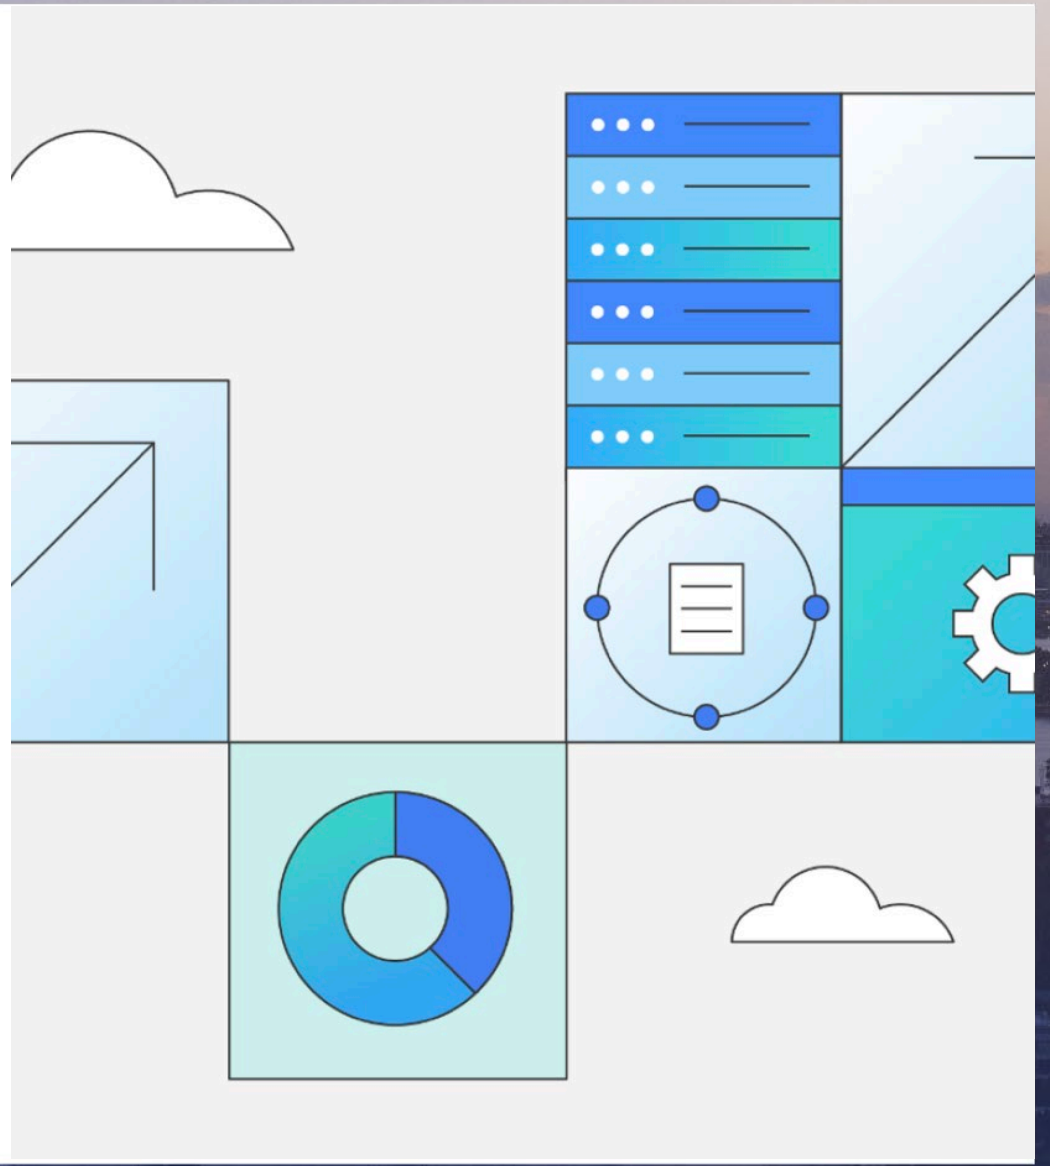

Four ways to ask Db2

AI_SIMILARITY

- Find similar or dissimilar entities

AI_SEMANTIC_CLUSTER

- Find entities that belong to a group of entities

AI_COMMONALITY

- Find most common entities, or outliers

AI_ANALOGY

- Analyze relationship between a pair of entities

Db2 for z/OS : SQL Data Insights

AI semantic functions in SQL

Find similar customers to account "abc" and rank top 20

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

 FETCH FIRST 20 ROWS ONLY ;
```

Re-discover Db2 data

- Industry's first vector DB for relational data
- Unsupervised deep learning - No data scientist required
- Everything is done on z/OS
- Utilize AI on-chip accelerator on Telum I (and II)
- Interface is your old "SQL", no-charged feature of Db2 13 for z/OS

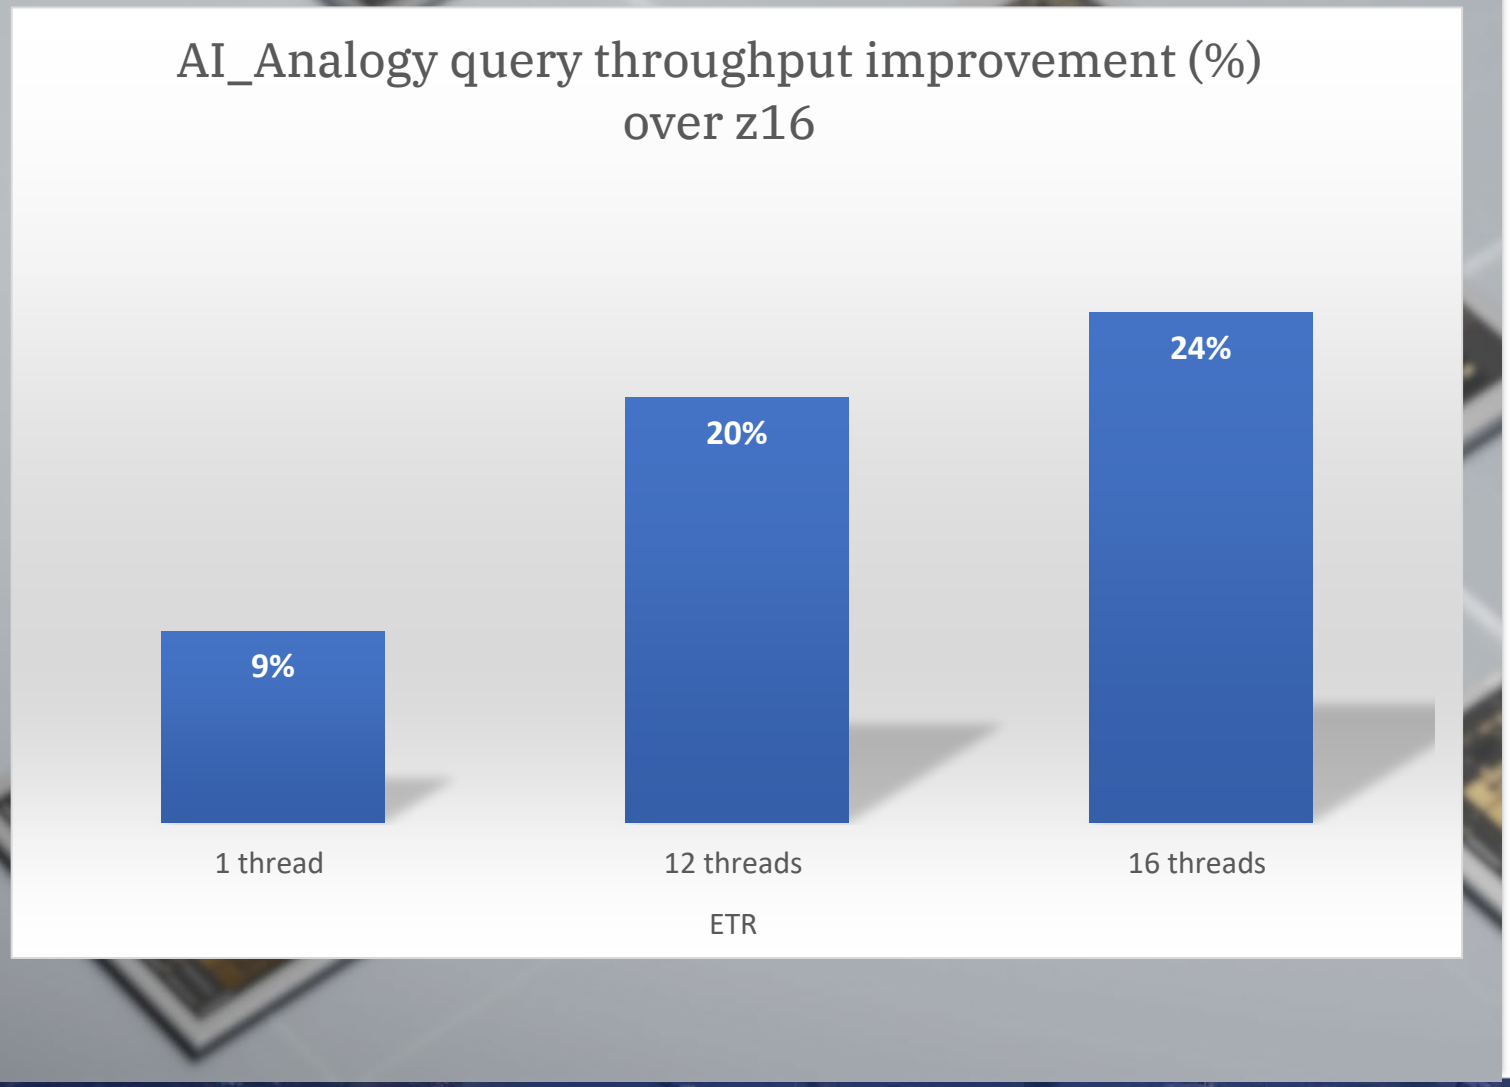

Telum II processor and SQL DI

- Better concurrency by intelligent routing within a drawer
  - 8 on-chip accelerators available within a drawer
  - Accelerators are shared within the drawers with all chips and cores connected
- Speed up SQL DI query execution

Db2 for z/OS : SQL Data Insights

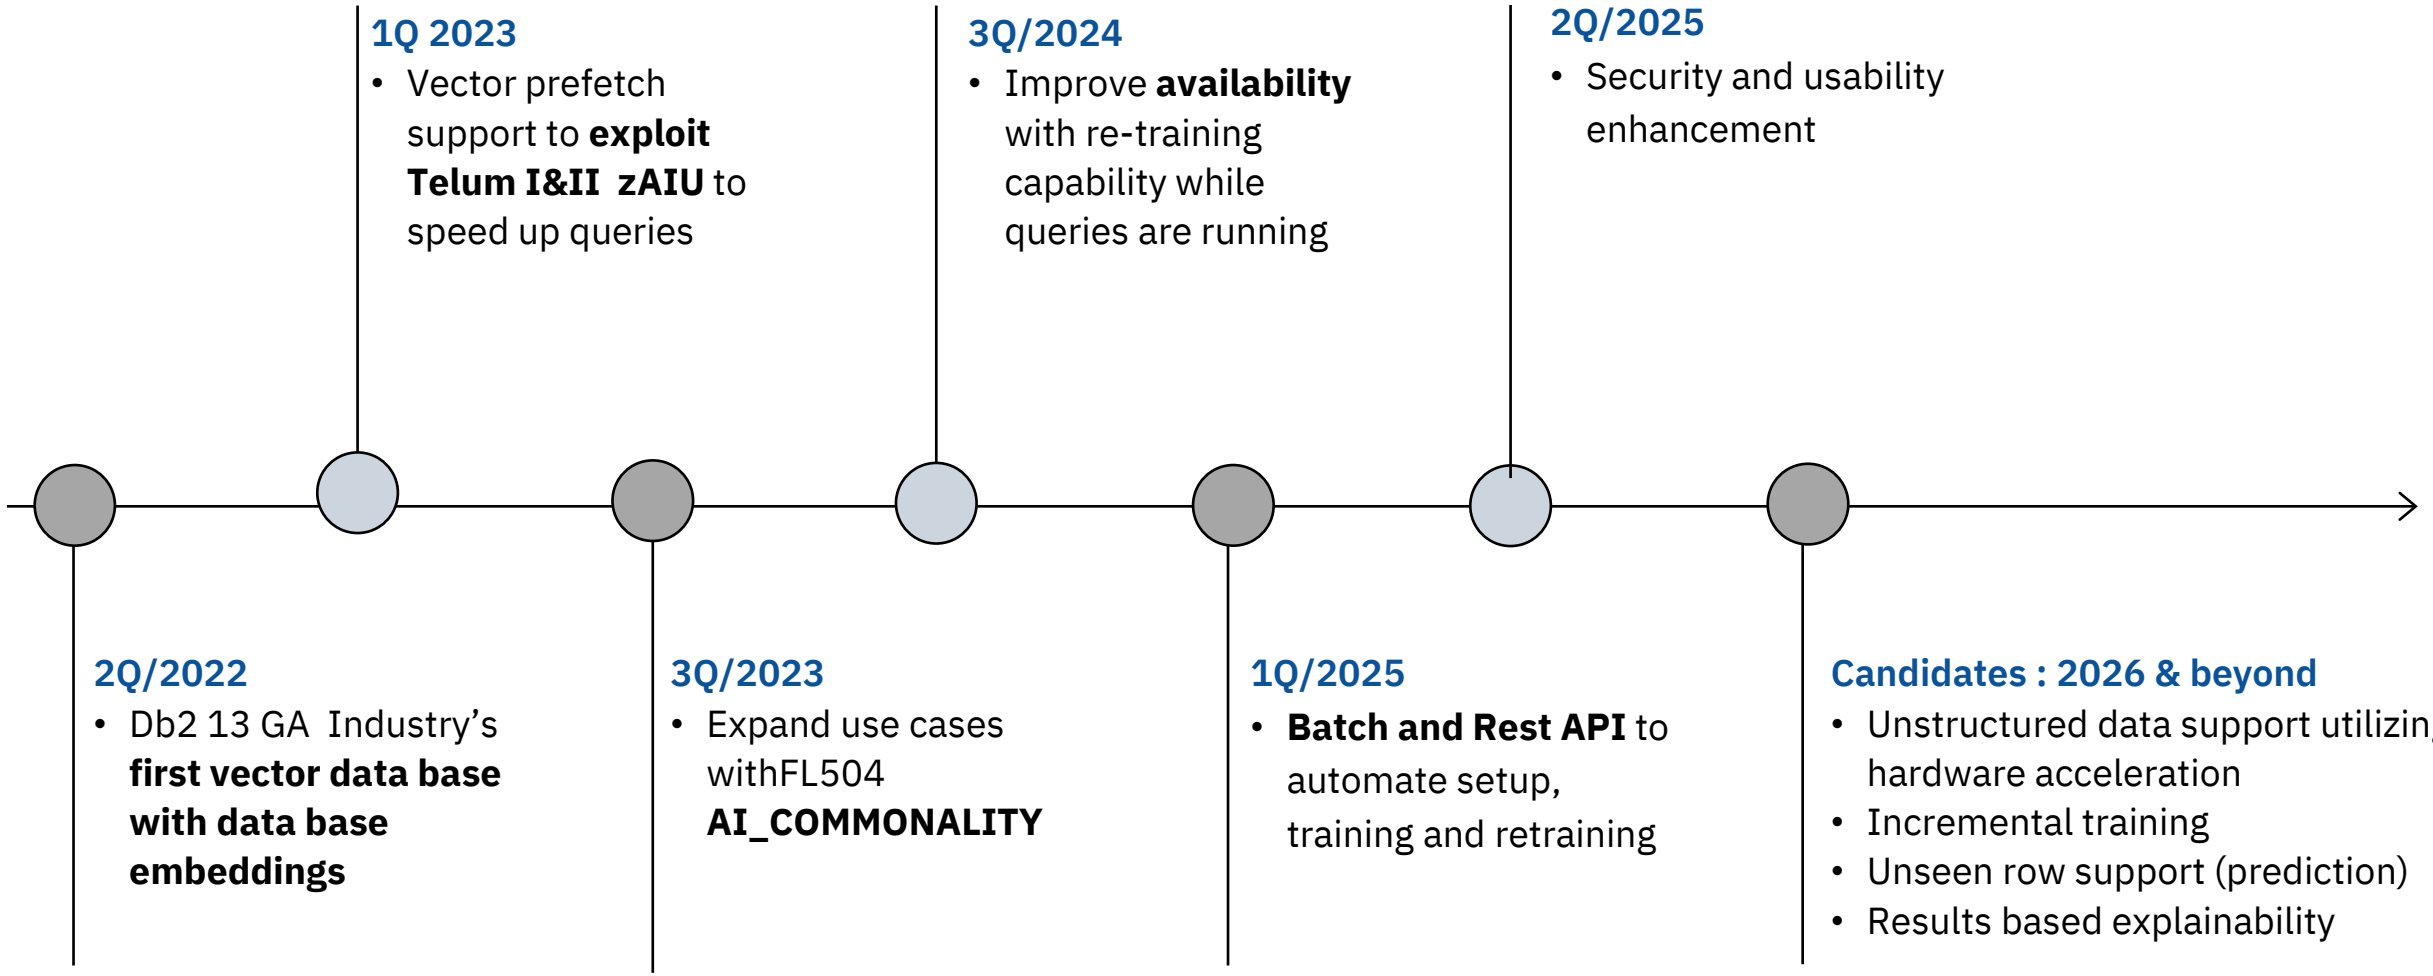

Supporting Unstructured Contents

Structured Content

Unstructured Content

ID	GENDER	AGE	CLASS	FLIGHT_DISTANCE	SEAT_COMFORT	GATE_LOCATION	LEG_ROOM_SERVICE	Departure Delay	Comments
795	Male	56	Eco	2456	0	5	5	0	Foods served were over-cooked and horrible !
796	Male	48	Eco Plus	1201	1	1	5	135	We were stuck in tarmac for 2 hours before told to get out. This occurs too often.
797	Female	8	Eco Plus	2364	0	2	0	0	Attendants were very nice (for a change) !
798	Female	28	Eco	2187	0	3	3	80	Delay caused me to miss the connection and caused me to pay for a hotel....

SQL DI (LDM)

Unified Database Embedding

LLM Embedding

AI_SIMILARITY(ID, 796)

Same query but including the insights from comments

SQL Data Insights : Potential Use Cases

Finance (Consumer Banking, Investment Advisors)

- Find customers with similar transactions
- Wealth Management:
  - Portfolio Recommendations
  - WM Advisor recommendations: Given an advisor, find an advisor similar or dissimilar to the given advisor
- Non-performing Asset Identification (NPA)
- Identify stocks with similar behavior as determined by their numerical attributes

Fraud detection

- Anti money laundering
- Account take-over

Insurance

- Claim analysis - Identify similar/dissimilar claims
- Underwriting – identify opportunities for revenue generation
- Evaluate risk profiles by analyzing patient profiles (e.g., symptoms, diagnosis...)

IoT

- Find households/hotel rooms with similar energy consumption patterns

Advanced sales prediction using external data

- Predict sales of new products to existing customer base

IT Analytics

- Identify tickets with similar behavior/patterns
- Identify customers with similar ticket patterns
- Performance Analysis over Numerical data such as evaluating system (e.g., Db2) execution performance using numerical attributes such as page faults, disk usage, cache misses, etc.

Customer analytics

- Find similar customers based on buying patterns
- Customer Churn Analytics

HR

- Find employees with similar skills and similar/different experience

Entity resolution/Data imputation for data quality

- Match rows that may have errors or NULL/missing values to fill or replace the values.
- Data with incorrect values can also be viewed as possible fraud (e.g., we looked at an auto Insurance claims dataset, where people were submitting fraudulent claims on the same car, simply by changing a few attributes in the claim: address, solicitor name, car information, etc.) By identifying possible claims as similar rows, we were able to identify possible fraudulent claims.

Db2 AI for z/OS – Db2ZAI

Improve Efficiency and Bring Db2 Operational Cost Down

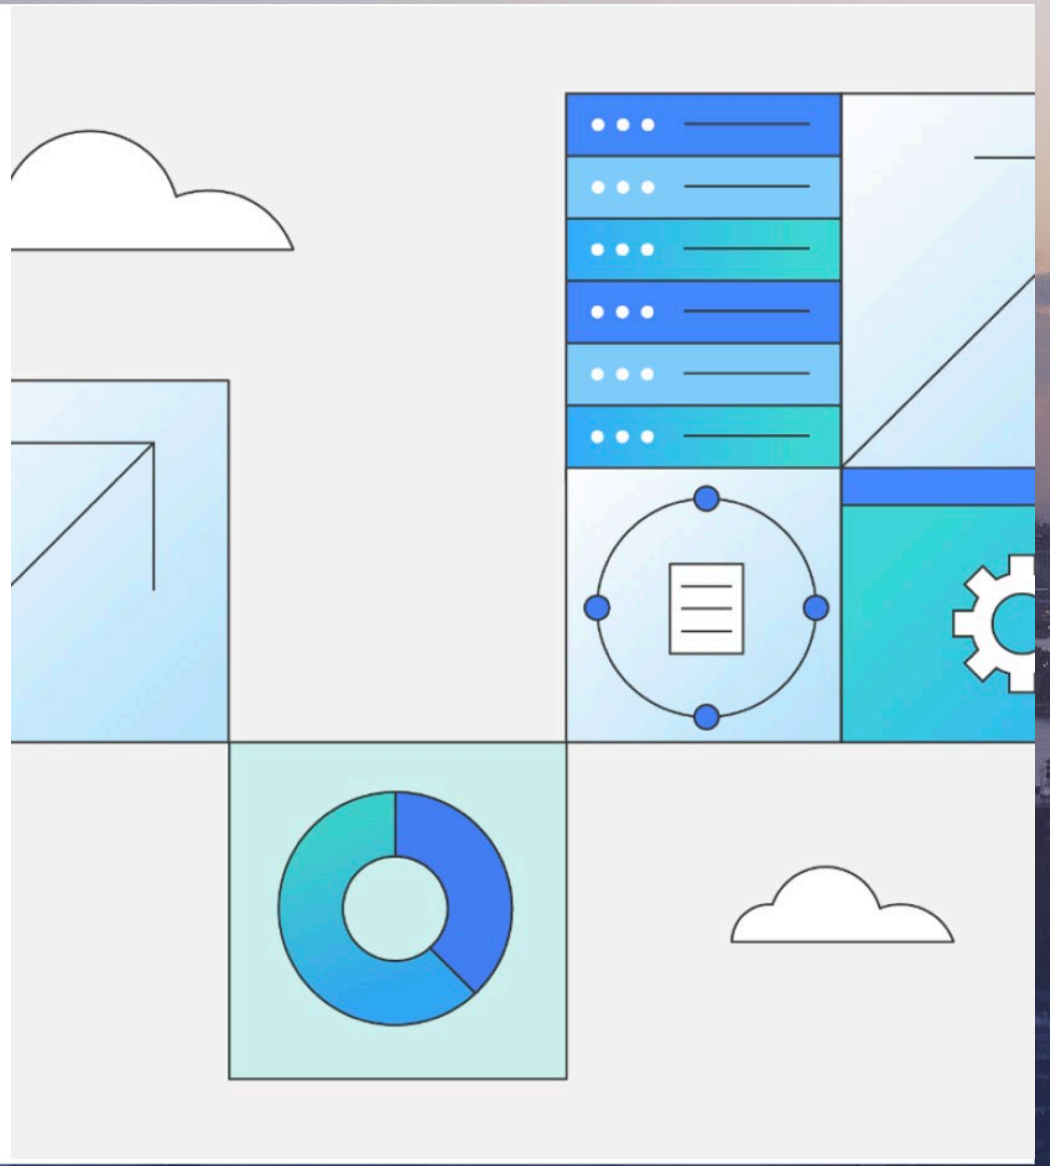

Db2 AI for z/OS

Collect the operational data and leverage AI to improve Db2 performance and resiliency

Data Collection

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

Learning and optimization

Db2ZAI learns, predicts and optimizes or recommends updates on Db2 execution behavior

SQL Optimization	**Reduce CPU consumption** CPU saving with better query access Path and stabilize production queries	
System Assessment	**Improve Productivity** Less time spent analyzing data, more on finding solutions	
Distributed Connection Control	**Protect critical resources** Act quickly, uncover issues early, prevent a flood of connections	

Bring operational AI to Db2 for z/OS to improve your database system health, efficiency and performance!

Db2ZAI – Journey to Self Tuning (1|2)

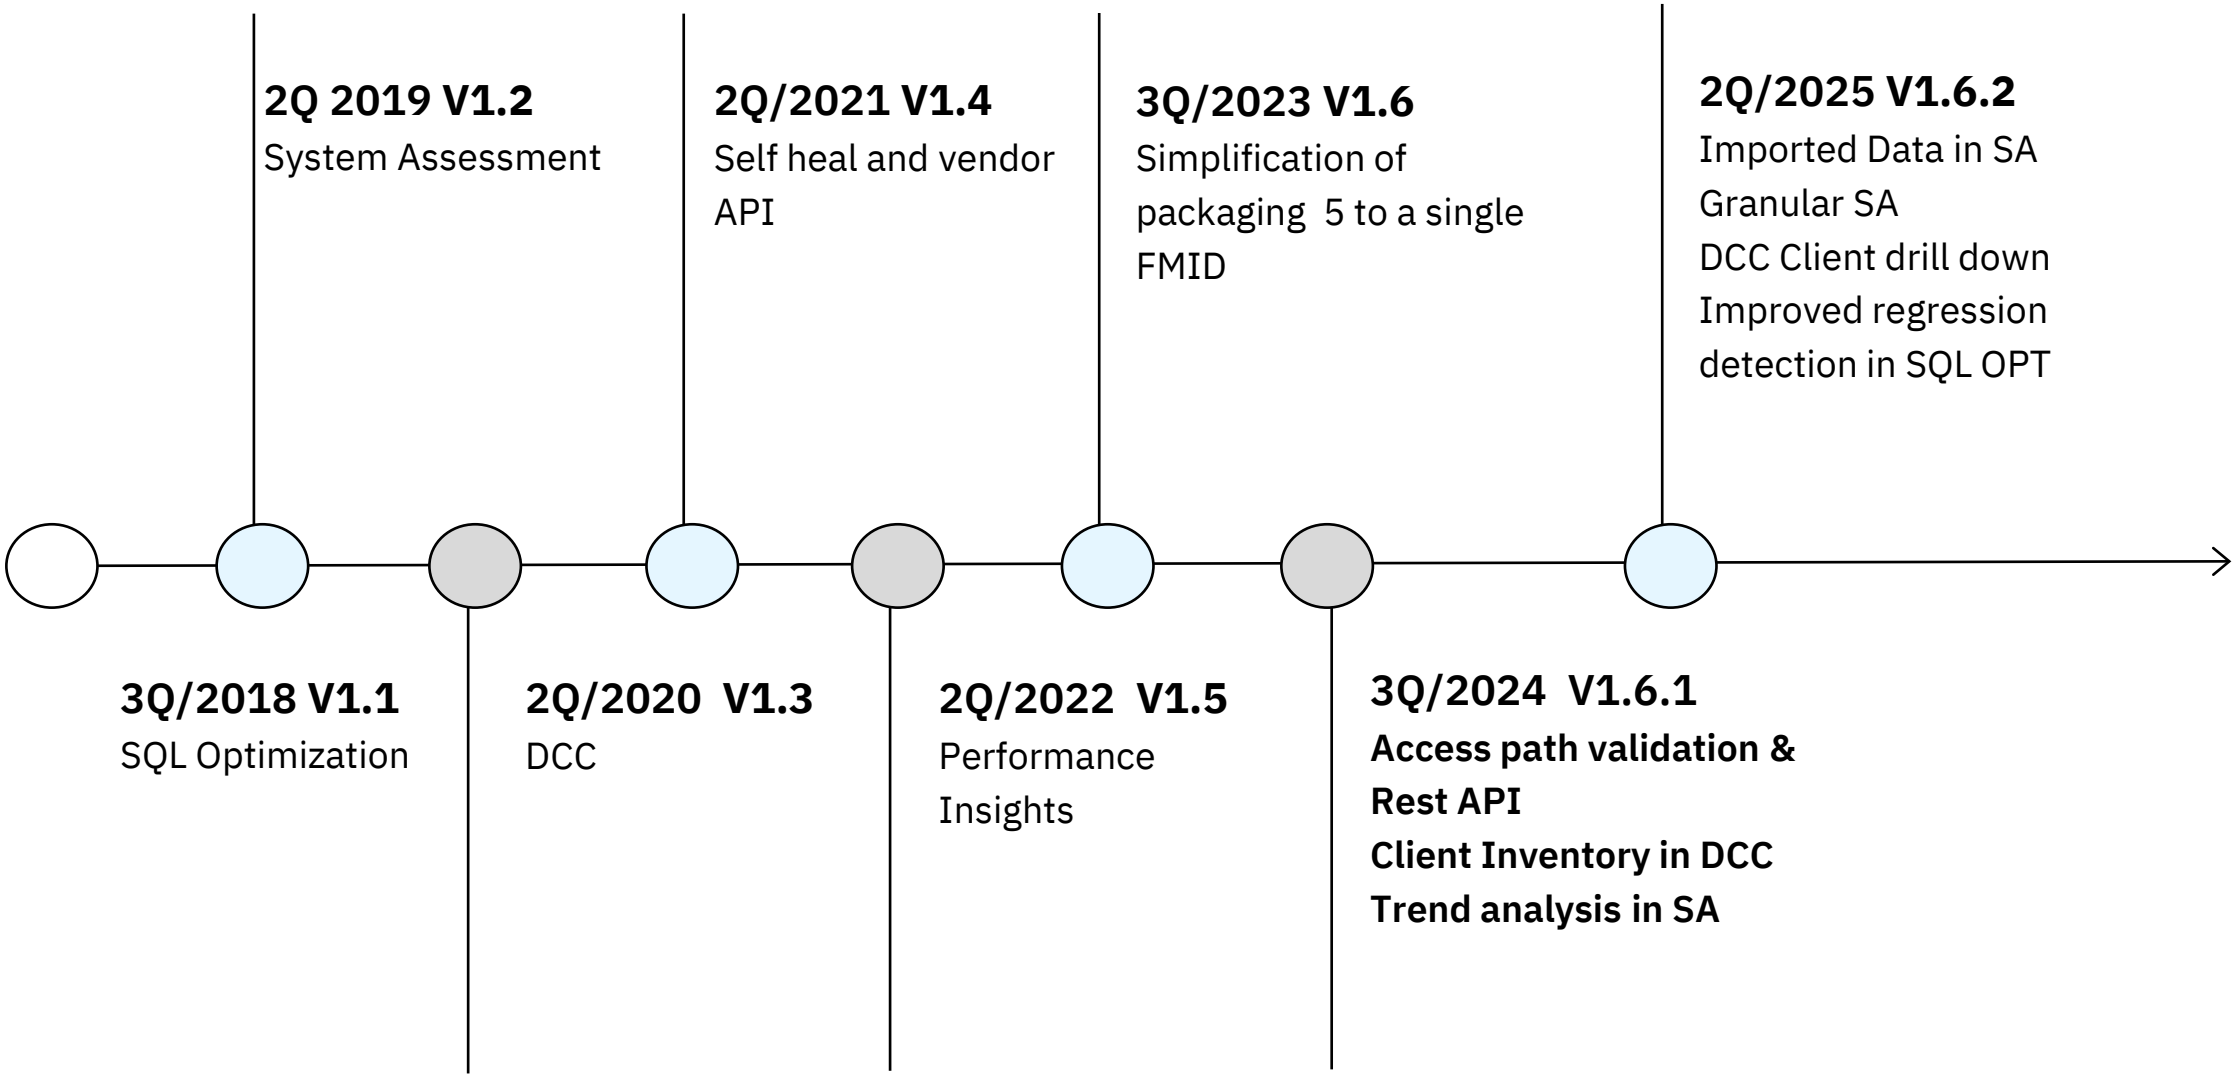

Db2ZAI – Journey to Self Tuning (2|2)

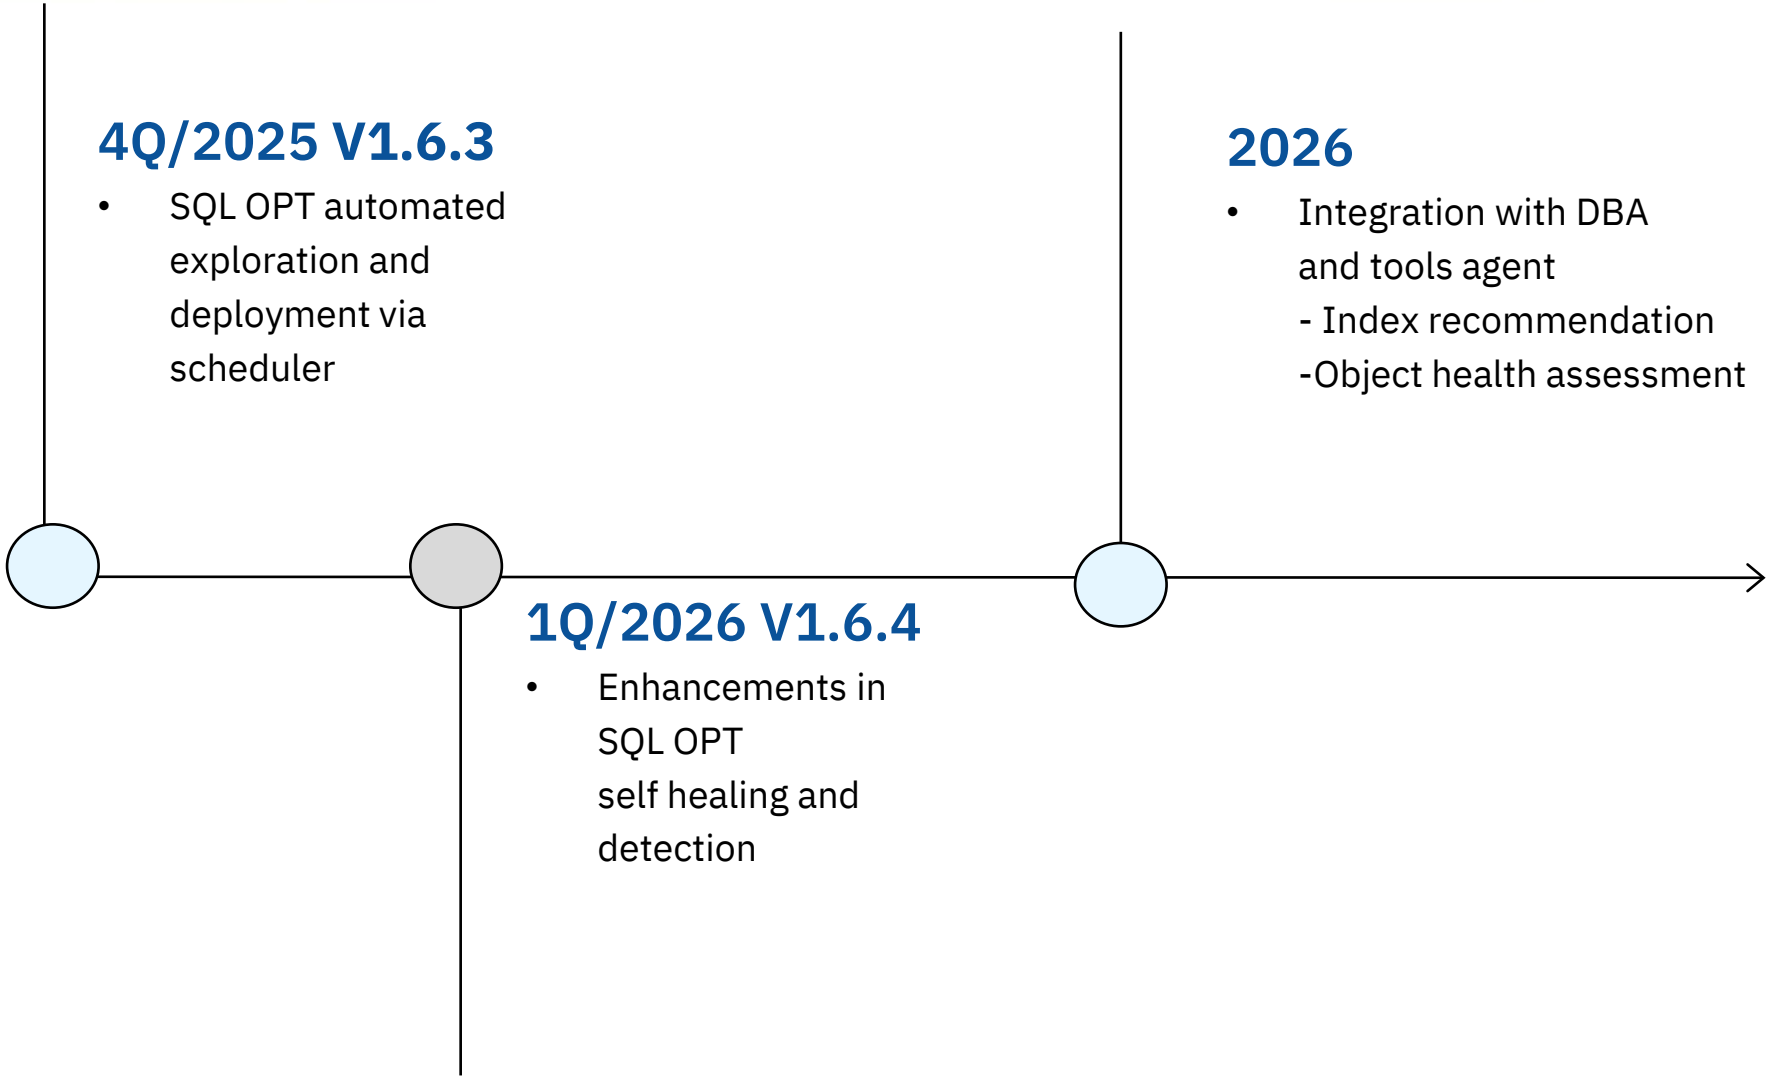

Index & Runstats Recommendation via Db2ZAI

- Starting with statement level recommendation
- What is the benefit over Query Tuning Services

1. Recommendation is based on the observation from execution information
   - e.g. filter factors and # of rows fetched
2. Consider the trade-off
   - Estimated benefit of adding index vs.
   - Cost of adding an index
3. Plan to expand to object level recommendation over multiple statements

Keep up with important Db2 for z/OS news …

IBM Db2 for z/OS and its ecosystem

February 2026

Customer Newsletter

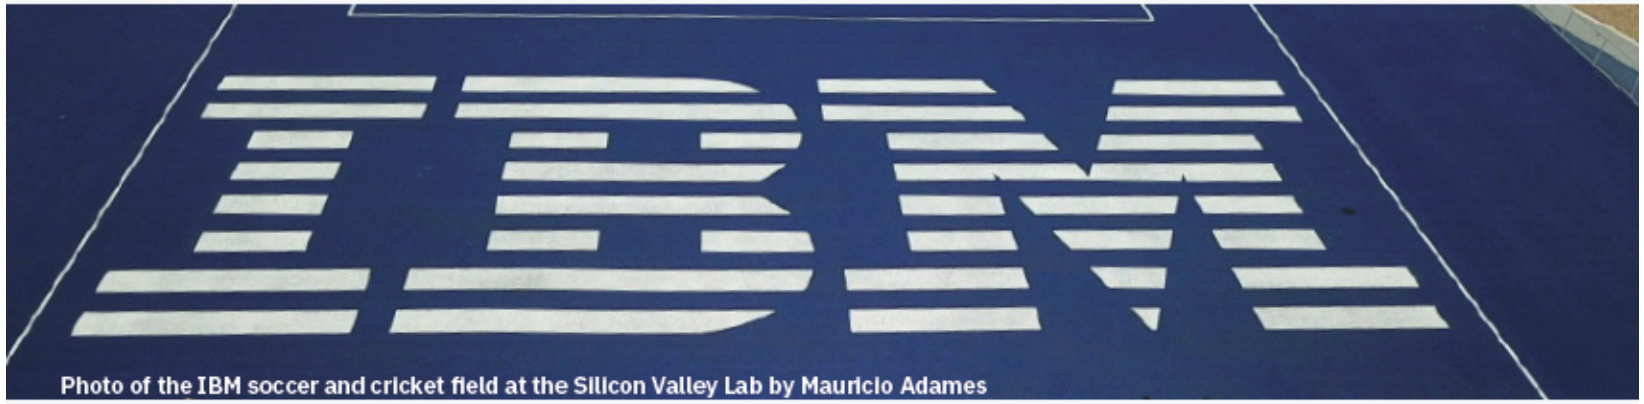

Photo of the IBM soccer and cricket field at the Silicon Valley Lab by Mauricio Adames

Welcome to the February edition of the Db2 for z/OS customer newsletter, your source for news and information about the Db2 for z/OS ecosystem, which includes:

- Db2 for z/OS
- Data Gate for watsonx
- Data Gate for z/OS
- Db2 AI for z/OS
- SQL Data Insights
- AI Optimizer for Z
- Db2 Analytics Accelerator
- Data Virtualization Manager for z/OS
- Db2 Developer Extension
- Db2 Administration Foundation
- Query Workload Tuner
- SQL Tuning Services
- The Db2 for z/OS tools

The newsletter is available on the Db2 for z/OS and its ecosystem page on the IBM Community site. Join the community and turn your notifications on to make sure you never miss the release of a new edition!

Please send your feedback and suggestions for this newsletter to Emily Alameida.

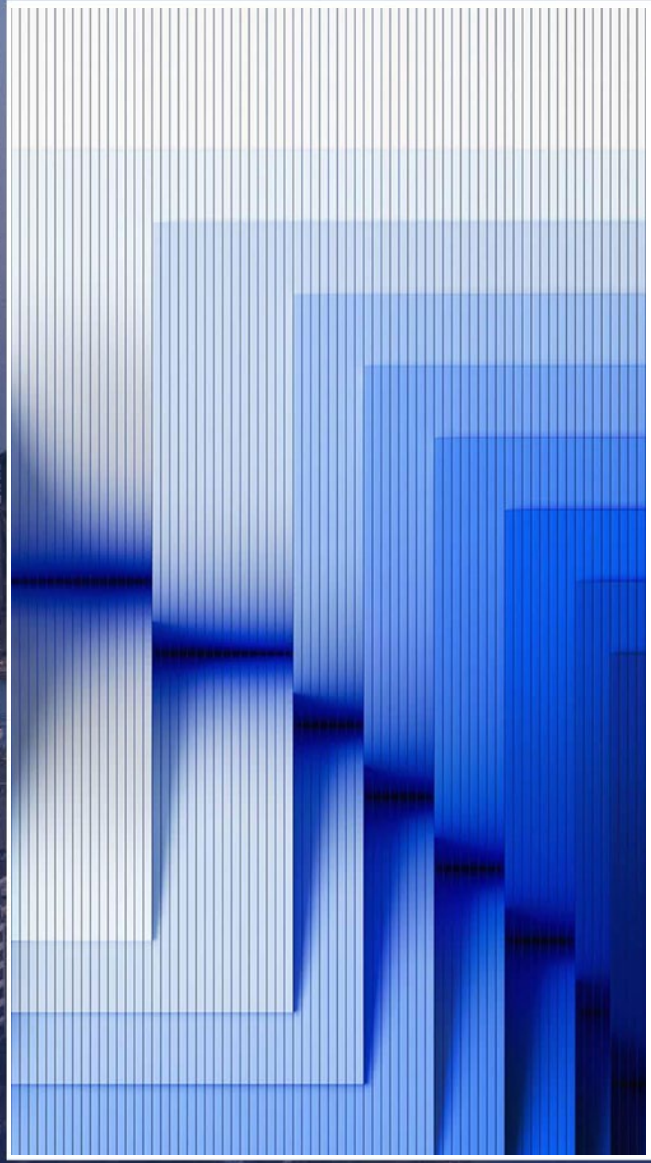

Summary

- Hybrid cloud data server, not data vault
- Unsurpassed investment in security, performance, resiliency, scalability, accessibility and 24x7 availability over four decades creating a strong foundation with a focus on the future
- Low-to-zero latency real-time insights from your most valuable data assets
- Simplified, optimized, automated administration and management of systems and applications.
  Utilize AI where it makes sense

IDUG

2026

Sydney | March 16 - 18

AU Db2 TECH CONFERENCE

Db2 for z/OS Trends & Directions

Sueli Almeida, IBM Silicon Valley Lab

Contact: sueli@us.ibm.com

Session Code: ####

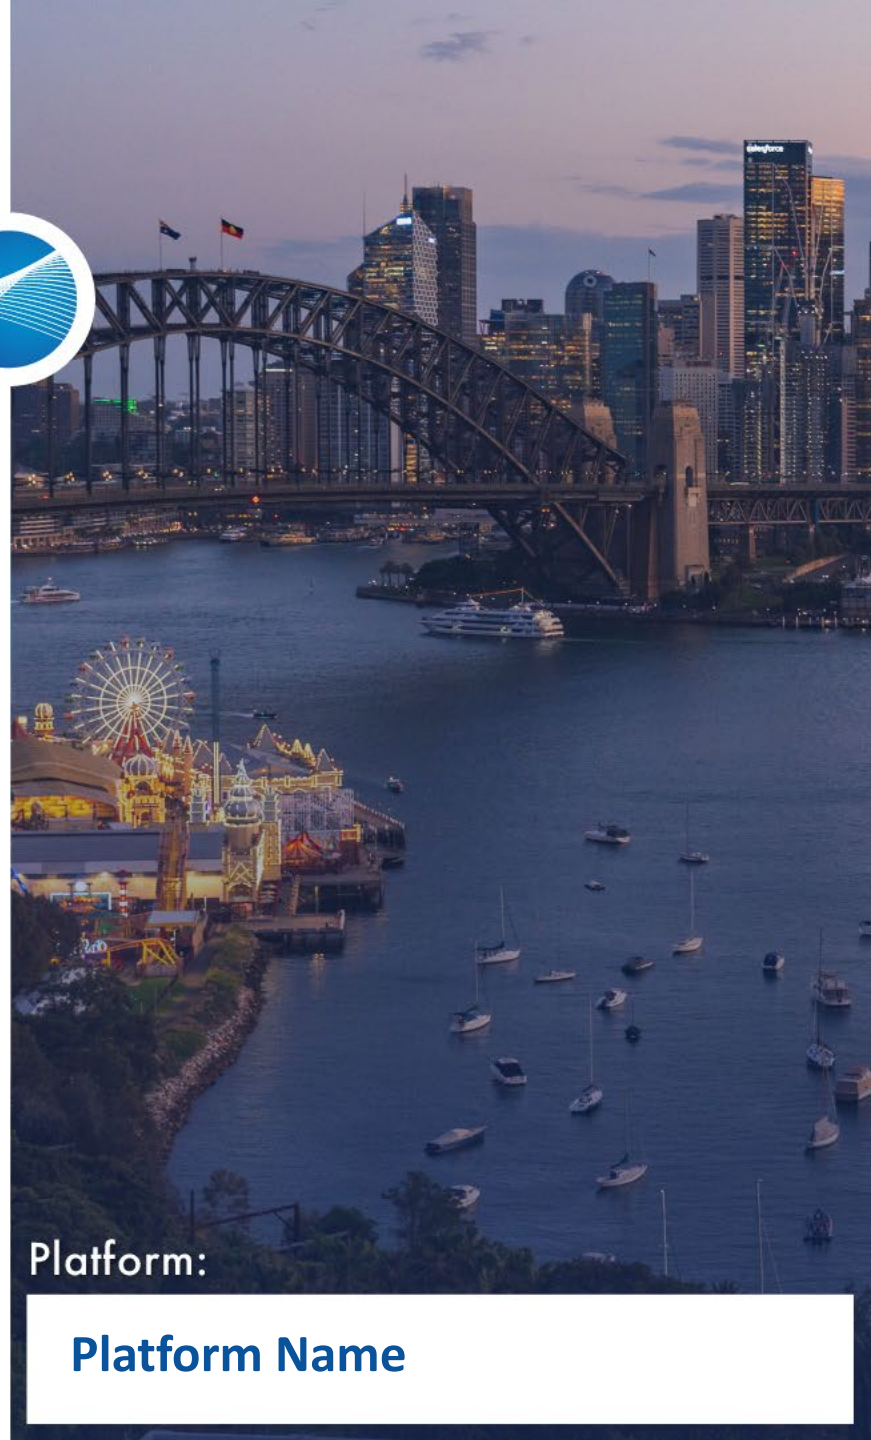


IDUG
2026 Australia Db2 Tech Conference