

IBM Db2

The Autonomous
Database Era has
Begun

Mike Springgay
Db2 Architect



Notices and disclaimers

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.

© 2026 International Business Machines Corporation.
All rights reserved.

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.

IBM is the
hybrid cloud +
AI company

Deploy Db2 anywhere your apps and dashboards run


Data proximity is critical for optimal performance. **Data locality** is critical for governance and access control. This is why we're **continuously expanding** where and how you can take advantage of **Db2**.

	Self-managed software	Cloud-managed IaaS	Cloud-managed container PaaS	Fully-managed SaaS	Bring-Your-Own-Cloud SaaS
Db2		   	  	 	
Db2 Warehouse		   	  	 	
Db2 pureScale		 			

Db2 & Db2 Warehouse as a service on Azure


(Bring Your Own Cloud)

Deploy **Db2 & Db2 Warehouse** as managed services on **Azure** in your own cloud account and Vnet



Managed Service

Managed by Db2 experts so customers can focus on what matters



Scalable & Elastic

Scalability through on-demand compute



Compliant & Secure

Take advantage of existing cloud security posture by deploying in the customer's cloud account




Connectivity

Connect to data sources directly in a customer's Virtual Private Cloud, no external networking needed

Fully managed Database Services

Customers focus on their Database applications and business, let our Database SaaS Offerings do the rest



Focus on...

- Schema design
- Query construction
- Query optimization
- Application modernization
- New cloud-native apps

Managed for you...

- Backup and recovery
- Patching
- Configuration
- Software upgrades
- Storage upgrades
- Server upgrades
- Hardware issues

Focus AI

Scale **AI** with **Db2** across the enterprise



AI for human productivity

Reduce management costs and improve database admin productivity

→ New **Db2 Genius Hub** with AI-powered **Database Assistant** for DBAs

→ New agentic AI capabilities delivering **massive productivity gains** via integrated agents in **Database Assistant**



AI for engine performance

Accelerate Db2 apps and analytics with built-in AI query optimization

→ Up to **3x faster** query performance with new **AI-based query optimizer**

→ **Enhancements & customization improvements** for AI query optimizer



AI for app builders

Combine curated data in Db2 with unstructured data to accelerate AI projects

→ **Build, train, evaluate & deploy** ML models with SQL, Python or R

→ **Vector store, index & similarity search** for RAG

→ **LangChain, LlamaIndex & MCP Server** for agentic AI



Integrations for AI


Integrations with watsonx platform makes it easy for customers to build AI

→ **watsonx.data:** share data with AI builders inside open data lakehouse

→ **watsonx.ai:** provide SQL access from Db2 directly to LLMs



Introducing the **Db2 autonomous database**, powered by the all new Db2 Genius Hub



Autonomous database that proactively monitors, detects, and reacts to issues, acting autonomously when appropriate, and engages human judgement only when necessary

→ **10x** Reduction
in database
management costs
with agentic
maintenance

→ **10x** Reduction
in manual
intervention
with agentic
healing

→ **10x** Reduction
in time to problem
resolution
with agentic
response

Genius runs on top of a network of AI agents that seamlessly **work together** to perform **activities** that power an **autonomous experience** for Db2 databases



Agentic Maintenance

10x reduction
in management costs

Automate critical tasks, such as tuning & patching. Each action is observable and reversible, ensuring operational autonomy without sacrificing control.

Activities

- **Engine Installation & Updates** **PREVIEW**
Deploy new or upgrade existing instances of Db2
- **Database Performance Optimization**
Stats refresh & reorgs to keep indexes and tables optimized
- **Health Check**
Proactive scans for common issues
- **Backups & Recovery Preparation** **COMING**
Automated, verified backups with audit trails
- **Integrity Checks & Cleanup** **COMING**
Validate constraints and maintain hygiene



Agentic Healing

10x reduction
in manual intervention

Continuously monitor usage patterns, detecting anomalies early. This reduces the likelihood of SEV-1 incidents and ensures planned work takes priority over firefighting.

Activities

- **Anomaly Detection & Resolution**
Identify unusual spikes in CPU, memory, storage, or queries
- **Locking Issue Resolution**
Detect blocking chains, deadlocks, and timeouts
- **Query Spilling Resolution**
Highlight when memory overflows to disk, explain impact, and recommend fixes
- **Workload Queueing Resolution**
Detect queue buildup, surface contributing workloads, and propose rebalancing actions



Agentic Response

10x reduction
in time to resolution

When issues arise, explain what changed, why it matters, and how to address it — saving hours of root-cause analysis. Every recommendation is auditable, helping build trust.

Activities

- **Performance Problem Determination**
Analyze performance issues across multiple diagnostic paths
- **What's Changed Analysis for Incident Management**
Identify schema, index, or config changes tied to incidents
- **Conversational Telemetry**
Natural language access to monitoring data and trends
- **Deep Research with Multi-Step Reasoning**
Answers from documentation, support run books & more
- **Incident Summary Report**
Generate root-cause reports and draft support tickets
- **Log Analysis for Advanced Problem Determination**
Track Db2 logs to quickly identify errors and issues

Hello, Admin

Key insights

core_banking_pr...

●

core_banking_...

●

retail_deposits_ua...

→

2 more

Last updated: 2 seconds ago

AI

667

Queries run

86%

Average CPU usage

17%

Average memory usage

40%

Peak CPU usage

20 Feb 2026 11AM

🔗

Select database tag

24 database tags

Select a database tag to view an overview of the associated databases.

Manage database tags

🏷️

All databases

🔴 5

50 databases

🟡 20

🏷️

Db2 Warehouse

🔴 1

7 databases

🟡 14

🏷️

Db2 Database

🔴 2

22 databases

🟡 5

🏷️

pureScale

🔴 1

21 clusters

🟡 5

+

➤

All databases

50 databases

🔍 Search databases

🏷️📄New connection +

Name	Tags	Alerts	CPU	Memory	Storage	Respon...
✓ core_banking_pr...	🏷️ prod us-west +1	🔴 1 🟡 2	30%	68%	74%	96 ms
✓ core_banking_dr...	🏷️ dr us-east +2	🔴 3 🟡 1	42%	73%	79%	47 ms
✓ retail_deposits_...	🏷️ uat retail +1	🔴 1 🟡 1	45%	77%	68%	51 ms
✓ loan_servicing_qa	🏷️ qa nonprod +5	🟡 4	33%	70%	61%	78 ms
✓ payments_hub_...	🏷️ sepa regulated +1	🟡 1	51%	65%	55%	45 ms

5 ▾

1–5 of 50 items

1 ▾ of 10 pages

◀▶

Notifications

Status: All notifications ▾ Sort by: Newest ▾

🟡 Rows read

core_banking_prd_us_east 01-01/23/26, 2:18:08 PM

🔴 Total activity time

core_banking_prd_us_east - January 23, 2026 2:15:05 PM

⋮

Hi I'm your Genius.
How can I help you today?

What's new
Db2 Genius is learning new skills!
Learn more →

- What can Db2 Genius do? ➤
- Retrieve schemas ➤
- How is my database performing? ➤
- Get query throughput ➤
- Analyze wait times for the past week ➤

core_banking_... ▾

Research Db2 topics ➤

Generate SQL

Deep research

Add connection

What can the **Genius** do for you?

Just ask...

You 11:51 PM

What else can you do?



11:51 PM

I'm your Db2 assistant, here to help with a wide range of database tasks:

Performance Troubleshooting & Analysis

Diagnose and resolve performance issues including slow queries, bottlenecks, lock contention, memory pressure, transaction log problems (log full, log pressure, blocking transactions), spilling, and queuing. I identify root causes and provide actionable recommendations.

Database State & Telemetry Retrieval

Retrieve current and historical database information including CPU usage, active queries, queued requests, application status, tablespace utilization, health checks, resilience status, configuration values, and system metrics.

Disk Space Management & Optimization

Analyze storage usage, identify space reclamation opportunities, recommend reorganization (reorg) candidates for tables and indexes, troubleshoot tablespace growth, and provide optimization strategies.

Db2 Documentation & Knowledge

Explain Db2 concepts, features, configuration parameters, SQL syntax, error codes, best practices, and internal behavior. Provide SQL queries and commands for you to use based on official documentation.

IBM Db2 Support Ticket Preparation

Help you prepare comprehensive support tickets for IBM Db2 support by gathering problem details, environment information, troubleshooting history, and recommending relevant diagnostics to collect.

Connection Information

You're currently connected to the **BLUDB** database as user **db2inst1** using the **PROD_ANALYTICS_SVL** profile.

What would you like help with?



PROD_ANALYTICS_SVL x | v

Type something...

SQL Generate SQL

Deep Research



IBM Db2 Bridge

One Data Movement
Approach for all Mission-
Critical Workloads

Launched
February 2025

IBM Db2 Bridge

Unified data movement solution for migrating database workloads, modernizing data infrastructure and for regular data movement.

70%

Reduction in migration cost with right strategy and tooling¹

Use Cases

Accommodates complex data movement

Schema-level selection of objects
Supports Row/Columnar organization, Little/Big Endian formats & all object types, security labels privileges and authorizations

Move workload with Db2 as target

IBM Db2
IBM Db2 Warehouse
IBM Sailfish IIAS (Source Only)
Power 10 Cloud Rack (Target Only)
Db2 Data files

Support all possible deployment options for Db2

On-prem to On-Prem
On-prem to Cloud
Cloud to On-Prem
Cloud to Cloud

Key Features



Data Integrity

Preserves data integrity by analyzing the dependencies between objects



OS agnostic

Db2 Bridge is not impacted by operating system differences such as endianness.



Multiple Db2 Versions

Support for Db2 versions as far back as 10.5 and as latest as 12.1



Moves all schema objects

Clones all tables, views, procedures, indexes and more



Fast unload with HPU

Ability to enable Optim High Performance Unload

¹ [Data Migration Strategies to Cut Down Migration Costs by 70%](#)

IBM Db2 Developer Extension

To learn, explore and thrive in Db2 !

IBM Db2 is the "gold standard" for backend reliability, powering

→ **90%**

of the world's top 100 banks

Yet, a Stack Overflow survey¹ reports

→ **1.8%**

of the developers report using it.

Let's build the gap and empower the developers with,

“IBM Db2 Developer Extension for VS Code”

... in making Db2 the database of choice for the developers, improving accessibility and ease of use.

Key Features



Create Db2 Instance

to quickly create a Db2 Community Edition instance so that developers can start running queries against a ready-to-use sample database and a user database just created.



Manage Connections

to create and manage connection profiles so that developers can efficiently connect to databases and run queries.



Explore Objects & Data

to efficiently explore Db2 objects and write embedded SQL so that developers can develop distributed applications more effectively.



Build and run SQL

to execute SQL Queries, browse through the results, and export the result sets



Syntax Checking

to write Db2 queries efficiently and effectively, reducing errors and improving overall development experience.

GA (planned)
Feb 19th, 2026

¹ <https://survey.stackoverflow.co/2024/technology>

Db2 Container Deployment



Db2 Container

Modernization of container with next generation operator

Make the service more intuitive and efficient: Improve user experience

Better align with Kubernetes/OpenShift-native deployment practices

Use Db2 interfaces wherever possible

Deliver same Db2 experience in cloud-native environments

Top 5 areas that the Next Gen design is targeting

1. Install & Deployment



2. Update/Upgrade



3. Security



4. Resiliency



5. Usability



Value-add Metrics: Kubernetes resource reduction

Category	Current Gen	Next Gen	Reduction
Secrets	13	3	↓ 77%
ConfigMaps	7	2	↓ 71%
Jobs	2	0	↓ 100%
Auxiliary Components	4	0	↓ 100%
Other	13	13	-
Overall	39	18	↓ 50%

Db2u Next Gen **Top 5 enhancements** coming in 2026

1. Install	2. Deploy	3. Resiliency	4. Security	5. Usability
Helm Install Support	Deployment Time (54% ↓) Image Footprint (86% ↓)	Brand new HA design Improve recovery from various failure events	CVE Reduction (89% ↓)	Environment Pre-Checker Maintenance mode annotation in CR for disruptive admin operations Integration with Db2 CSB Model

Db2 Software

Db2 Delivery Vehicles (simplified)

Early Access Program (EAP)

- Delivered every month for current development branch
- Provides early access to new development features.
- Predefined Vehicle to validate features early.

New Function

- Delivered through modification (“mod”) Packs
- Identified by a new point modification identifier
- e.g. 12.1.1

Service Stream (aka Fixes)

- Delivered through cumulative special build (CSB) process
- Available on demand when customer requires immediate fix for a KI
- Fully validated CSBs published every 4-8 weeks along with list of KIs fixed. Security fixes delivered through CSB stream.

Db2 12 Release schedule



Db2 12 Release schedule



Scale **AI** with **Db2** across the enterprise



AI for human productivity

Reduce management costs and improve database admin productivity

→ New **Db2 Genius Hub** with AI-powered **Database Assistant** for DBAs

→ New agentic AI capabilities delivering **massive productivity gains** via integrated agents in **Database Assistant**



AI for engine performance

Accelerate Db2 apps and analytics with built-in AI query optimization

→ Up to **3x faster** query performance with new **AI-based query optimizer**

→ **Enhancements & customization improvements** for AI query optimizer



AI for app builders

Combine curated data in Db2 with unstructured data to accelerate AI projects

→ **Build, train, evaluate & deploy** ML models with SQL, Python or R

→ **Vector store, index & similarity search** for RAG

→ **LangChain, LlamaIndex & MCP Server** for agentic AI



Integrations for AI

Integrations with watsonx platform makes it easy for customers to build AI

→ **watsonx.data:** share data with AI builders inside open data lakehouse

→ **watsonx.ai:** provide SQL access from Db2 directly to LLMs

Db2 12.1.0

GA November 2024

200+ new features & enhancements across Db2, Db2 pureScale & Db2 Warehouse

Db2

- **AI-powered query optimizer**, allows Db2 to continuously learn from customer's queries and achieve up to 3x query performance improvement
- **Improvements to backup performance** with multiple threads per table space
- **Apple Silicon driver support** for developers on macOS
- **Security enhancements** with AUDIT exceptions, Trusted Context and data masking
- **Recovery time improvements** in the unlikely event of crash
- **Federation updates** to Snowflake, Oracle 23ai, MariaDB, MongoDB and SAP with across-the-board performance improvements

Db2 Warehouse


- **Name space separation** with TENANT construct, creating a logical separation between one or more database schemas
- **Cloud object storage performance & scalability enhancements** for large deployments
- **Schema evolution with DROP and RENAME** support for online schema updates to columnar tables
- **UPDATE and JOIN** performance enhancements for columnar tables
- **Logical backup/restore** experience improvements
- **Enhancements to online columnar table move** with ADMIN_MOVE_TABLE

Db2 pureScale

- **SSL encryption** support for Db2 HADR configurations
- **Cluster recovery** through offline drop member
- **TSA replaced with Pacemaker on Linux** technology for cluster management, leading to faster failure recovery times

Backup Performance – Intra-tablespace Parallelism


- Single node with 1FC Card:
- 2.7X faster in extreme skewed case, best scenario
- Flat for balanced case
- Improvement limited by headroom of IO bandwidth on the system



AI-powered query optimizer

Db2 12.1 integrates advanced AI optimizer capabilities, allowing businesses to harness the power of AI directly within their database their operations. This elevates query performance through automated performance tuning, predictive analytics, and intelligent query optimization.

AI allows Db2 to continuously learn from client queries, automate SQL tuning with AI, and optimize performance 3X faster.



Benefits




3X benefit in scenarios simulated in-house




In practice the average benefit will depend on workload



Minimize errors in cardinality estimates produces reliable performance



0 indicates perfect estimate



Db2 12.1.1

GA February 2025

Minor update that builds on cloud and query optimization capabilities released in 12.1.0

Db2

- **Security enhancements across the board**

Db2 Warehouse

- **Improvements to AI-based query optimizer**, leading to more accurate and faster results
- **Db2 Warehouse on Linux on Power (PPCLE)** now supports open data and table formats for data lake integration

Db2 pureScale

- **Db2 pureScale available as a self-managed deployment on Azure**, accessible through the Azure Marketplace and deployable within minutes
- **Online index reorg** for Db2 pureScale is now available, enabling read and write access to the underlying table during maintenance
- **Support for fast I/O fencing for NVMe shared storage** is now available for Db2 pureScale self-managed deployment on AWS

Scale **AI** with **Db2** across the enterprise



AI for human productivity

Reduce management costs and improve database admin productivity

→ New **Db2 Genius Hub** with AI-powered **Database Assistant** for DBAs

→ New agentic AI capabilities delivering **massive productivity gains** via integrated agents in **Database Assistant**



AI for engine performance

Accelerate Db2 apps and analytics with built-in AI query optimization

→ Up to **3x faster** query performance with new **AI-based query optimizer**

→ **Enhancements & customization improvements** for AI query optimizer



AI for app builders

Combine curated data in Db2 with unstructured data to accelerate AI projects

→ **Build, train, evaluate & deploy** ML models with SQL, Python or R

→ **Vector store, index & similarity search** for RAG

→ **LangChain, LlamaIndex & MCP Server** for agentic AI



Integrations for AI

Integrations with watsonx platform makes it easy for customers to build AI

→ **watsonx.data:** share data with AI builders inside open data lakehouse

→ **watsonx.ai:** provide SQL access from Db2 directly to LLMs

Db2 12.1.2

GA June 2025

Packed with new features designed to unlock the full potential of your data and drive the next generation of AI and cloud applications.

Db2

- **AI features** with **vector store** and **similarity search** for advanced AI applications
- **Backup performance improvement** with reduced flushing for backups by default
- **Federation enhancement** with Apache Cassandra support via ODBC
- **DATAMASK on READ** to obfuscate sensitive fields
- **TLS encryption monitoring enhancement** to track the impact of encryption on performance
- **Option to restrict world-writable file permissions** for better access control

Db2 pureScale

- **Optimized performance** with low latency network support powered by Elastic Fabric Adapter (EFA) on AWS
- **Socket Support in zLinux** brings cost-efficient configuration for non-production workload
- **z/VM Support** for additional deployment options on IBM Z

Db2 Warehouse

- **Azure BLOB Storage Support** for table storage, significantly reducing the cost of storing data without sacrificing performance
- **UPDATE and DELETE** support for Iceberg tables

Db2 12.1.2

Vector store + similarity search



Real-time **fraudulent detection** for financial transactions



Real-time **image matching** of people / places / objects



Real-time **customer profile matching** for risk and creditworthiness



Retail product **recommendations** based on prior purchases or preferences



Semantic **document search** (chat with your docs), including RAG

```
CREATE TABLE VECTOR_EMBEDDINGS (  
    vector1 VECTOR(100, REAL),  
    vector2 VECTOR(300, INT8)  
)
```

VECTOR

Constructor for the VECTOR type

Example: VECTOR('[-0.0762727, 0.0559632]', REAL)

VECTOR_DISTANCE

Similarity search for two or more vectors

Example: VECTOR_DISTANCE(vector1, vector2, COSINE)

VECTOR_SERIALIZE

Serialize vector type to string

Example: VECTOR_SERIALIZE(vector1)

VECTOR_DIMENSION_COUNT

Count the elements in a vector

Example: VECTOR_DIMENSION_COUNT(vector1)

Db2 12.1.3

November 2025

New AI capabilities to supercharge your AI initiatives across the enterprise

Db2

- **Connectors for Python frameworks LangChain and LlamaIndex** to support development of LLM-driven apps
- **Vector support extended** to utilities, external tables and routines
- **Online table reorg of range partitioned tables** with global indexes for seamless maintenance
- **Generate UUID** scalar function
- **DBMS_LOCK** transaction scoped application lock support
- **Db2diag** tool optionally output to JSON
- **Audit to COS** support writing in text (CSB) format

Db2 Warehouse

- **AI Optimizer updates** join cardinality estimation
- **Access Plan Stability** with ability to auto-generate optimization profiles for a query using the Explain facility
- **Flexibility to disable row modification tracking for logical backups** using CREATE SCHEMA and ALTER SCHEMA
- **ALTER the precision for decimal columns** in column-organized tables
- **Columnar Group-by** performance and improved memory handling


Db2 pureScale

- **HADR support in pureScale environments** with mixed topology between primary and standby clusters
- **Upgrade or restore database** using a down-level backup image with different topology
- **Quiesce tablespaces** to make them read-only in pureScale environments using db2dart
- **Performance enhancements to AWS EFA (up to 40%)**
- **Support for GIPC clusters with Pacemaker** to enable multi-AZ active-active clusters in cloud environments like AWS and Azure


Db2 12.1.3

AI Query Optimizer

Default 99 DS Queries



"Join Intensive" Modified DS Queries



TPC-DS Workload Configuration

- **Scale Factor:** 3TB. TPC-DS is a highly studied and tuned workload.
- **Traditional Queries:** normal 99 TPC-DS query set
- **Modified Queries:** based on original queries with more joins, different driving fact tables, and more number of tuples being processed by certain key joins

Scale **AI** with **Db2** across the enterprise



AI for human productivity

Reduce management costs and improve database admin productivity

→ New **Db2 Genius Hub** with AI-powered **Database Assistant** for DBAs

→ New agentic AI capabilities delivering **massive productivity gains** via integrated agents in **Database Assistant**



AI for engine performance

Accelerate Db2 apps and analytics with built-in AI query optimization

→ Up to **3x faster** query performance with new **AI-based query optimizer**

→ **Enhancements & customization improvements** for AI query optimizer



AI for app builders

Combine curated data in Db2 with unstructured data to accelerate AI projects

→ **Build, train, evaluate & deploy** ML models with SQL, Python or R

→ **Vector store, index & similarity search** for RAG

→ **LangChain, LlamaIndex & MCP Server** for agentic AI



Integrations for AI

Integrations with watsonx platform makes it easy for customers to build AI

→ **watsonx.data:** share data with AI builders inside open data lakehouse

→ **watsonx.ai:** provide SQL access from Db2 directly to LLMs

Db2 12.1.4

March 2026

Minor update with upgrades to technology stack for enhanced security and performance for modern workloads

Db2

- **Red Hat Enterprise Linux 10** in both standalone and DPF configurations.
- **DataStax federation** via JDBC and ODBC wrappers
- **Extract audit events from logs in cloud object storage** to support centralized retention and compliance workflows.
- **Asynchronous history file pruning** with online backup

Db2 Warehouse

- **ALTER the scale for decimal columns** in column-organized tables
- **ALTER SMALLINT to INT** in column-organized tables
- **Datalake Table (ODF) support** on Db2 12.1 non-containerized deployments for DPF

Db2 pureScale

- **Technology stack updates**

12.1.4 Data Lake Tables

Warehouse Augmentation

Gain new insights from your warehouse data by combining Db2 Warehouse and data lakes platform data through open formats engine.

Optimize Workloads

Use the most appropriate tool for the task at hand without having to move or copy the data

Share data through an open format

Eliminate data silos by sharing Db2 tables with data lakes and lakehouse engines.

12.1.4 Data Lake Tables

Located in Object Storage

- S3 Compatible (IBM COS, AWS S3)
- Leverages Remote Storage Aliases to provide credentials and bucket

Used to access open data formats

- Parquet, ORC, Avro
- Limited support JSON and TEXT

Used to access open table format ICEBERG

Current Availability

- Containers and Db2 Warehouse SaaS
- Traditional install Db2 DPF installations on x86
 - Single Database per Instance can be enable for data lake
 - Database must be UTF-8 with IDENTITY Collation

ICEBERG Tables

- CREATE (DDL or GUI)
 - Includes partitioning support
- SELECT (aka READ)
- INSERT
- DELETE
- UPDATE
- Time Travel
 - Snapshot Rollback
 - Snapshot Management

REST Catalog for DATALAKE Tables

Db2 12.1.5

June 2026

New AI capabilities to supercharge your AI initiatives across the enterprise

Db2

- **Vector Indexing with DiskANN integration** for fast access to 1 billion vectors
- **LLM integration with watsonx.ai and OpenAI models** to generate vector embeddings for data and text with SQL statements
- **Support more than 3 HADR standbys**
- **Cloud Object Storage** Store and access backups and load copies on S3 from AIX
- Flexibility to **restrict file permissions for world-readable and world-executable files**
- **Upgrade to GSKit 9** for FIPS140-3 and more encryption updates
- **Enable automatic archive log pruning with non-Db2 snapshot backups**
- **REORG** Improve reorg check F2 for range partition tables


Db2 Warehouse

- **Enable SET NOT NULL** in column-organized tables
- **Enable ALTER SMALLINT/INT/BIGINT** data types
- **AI Optimizer enhancements** for Cardinality estimation, including join card string support, and 2-dimension support
- **Quoted values supported when unloading External Tables**

Db2 pureScale

- **Add disk as Tiebreaker mechanism to Pacemaker**
- **Fencing supported with quorum device** in automated HADR with Pacemaker

Build LLM Apps with Db2




IBM Community

Join the IBM TechXchange Community Today!



Engage Directly with IBM
Champions and subject matter
experts! Learn tips, tricks, and best
practices from your industry peers!



Understand the fundamental
concepts, models, tools, and
applications of generative AI





[Join Now!](#)



Participate in community
challenges, ongoing learning, and
certification opportunities.
Receive **an IBM Digital Badge!**



Register for online webinars and
in-person events!



Call for Sessions

Submission deadline: [May 22](#)

IBM TechXchange 2026

October 26–29 | Atlanta, GA

Clients, Partners & IBMers are eligible to submit sessions

Share your expertise with thousands of technical professionals!

What we are looking for

- Real-life use cases and implementation stories
- Technical deep dives (non-promotional)
- Demonstrated expertise with IBM products
- Sessions aligned to 2026 priorities: enterprise AI applications, key products, technical skills
- A compelling title and concise, clear abstract

Session types



Technology Breakout (45 minutes)

In-depth technical session featuring architectures, demos, best practices, client/partner stories



Tech Talk (20 minutes)

Fast-paced spotlight on use cases, technical trends, and special-topic insights.

Client & partner submissions only for Tech Talks.

Technology areas

AI Applications Data Management
Cloud & Platform Engineering Infrastructure
Security, Risk & Compliance

Special perks

- A complimentary full conference pass (that's a \$1599 USD value!)
- Speaker training and content support
- Social tools to promote your session

Submission deadline: [May 22](#)

Get started



