


Db2: What's new and What's next

Mike Springgay

IBM, Db2 Chief Architect



Notices and disclaimers

- © 2025 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.
- 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.

IBM Db2: 30+ years of innovation

1970s

SQL invented by
Edgar F. Codd at IBM

1983

Db2 for Z/OS is born
(40+ years ago)

1993

Db2 LUW is born
(30+ years ago)

Db2/2 V1



1993

Db2/6000 V1



1994

Db2 PE



1994

Db2 2.1



1995-1996

Db2 5-6



1997-1999

Db2 7-8



2001-2004

Db2 9-9.8



2006-2010

Db2 10.1



2012

Db2 10.5



2013

Db2 11.1



2016

Db2 11.5



2019

Db2 11.5.4



2020

Db2 11.5.5



2020

Db2 11.5.6



2021

Db2 11.5.7



2021

Db2 11.5.8



2022

Db2 11.5.9



2023

Db2 12.1.0



2024

- HADR
- .NET, JDBS, SQLJ, OLE drivers
- VARXXX, XLOBs
- OS/2, AIX, Windows, Linux, Solaris, HPUX
- Text Analytics
- Shared-Nothing Scale-out for OLAP (DPF)
- Granular backups
- Spatial Analytics
- Sequences
- Query Patroller (WLM)
- Db2 Connect
- Data Links
- Data Joiner (Federation)
- Connection pooling
- LDAP integration
- Unicode
- AST/MQT
- Mobile Satellite
- Triggers
- Appliances
- Shared-disk, scale-out for OLTP (pureScale)
- PHP, Perl, Python, RoR, ADO, PL/SQL
- Label-based access control
- Row/column access control
- Roles
- Range partitioning
- pureXML (NoSQL)
- Autonomics
- Multi-tiered storage
- Native encryption
- Audit
- Trusted context
- Multi-dimensional clustering (MDC)
- JSON/BSON support
- Oracle application compatibility
- Compression (tables, index, temp tables)
- Continual data ingestion
- Native OLAP functions
- Native WLM
- Online utilities
- Columnar (BLU)
- Temporal tables
- PostgreSQL compatibility (for NZ workloads)
- Db2 on Cloud (DBaaS)
- External tables
- Event processing
- ML optimizer
- In-database ML
- Advanced log space management
- Graph
- Data virtualization
- Red Hat OpenShift support
- Schema-level security
- Schema-level recovery
- Adaptive Workload Management
- REST APIs
- Namespace Separation (tenancy)

Db2 Four Big Bets for 2024

Db2 on Cloud Hyperscalers

AWS RDS for Db2 licensing and feature enhancements

Db2 on Cloud Next Gen in IBM Cloud

pureScale IaaS in AWS and Azure

Db2 Warehouse Gen3 on IBM Cloud

Fully managed cloud data warehouse featuring Db2 tables on Cloud Object Storage, support for open data formats and watsonx integration

Matches AWS offering

Db2 infused with Generative AI

We're adding Gen AI capabilities into Db2.

Db2 12.1

Full software release targeting all platforms and deployment options

Major release providing new base and jumping off point for continues mod pack deliveries throughout 2025

IBM Db2 + Amazon Web Services

Partnering closely with Amazon to bring our Db2 offerings to AWS

Other offerings available, including:

- **Db2 RDS for OLTP workloads (managed service)**
- **Db2 Warehouse on Cloud (managed service)**
- **Db2 pureScale on AWS**
- **Db2 Container reference architecture**

IBM Signs Strategic Collaboration Agreement with Amazon Web Services to Deliver IBM Software as-a-Service on AWS

- Building on IBM Software being available as-a-Service on IBM Cloud, this first-of-its-kind agreement between IBM and AWS will provide clients with access to IBM Software that runs cloud-native on AWS

May 11, 2022



ARMONK, N.Y., May 11, 2022 /PRNewswire/ -- IBM (NYSE: IBM) today announced that it has signed a Strategic Collaboration Agreement (SCA) with Amazon Web Services, Inc. (AWS), with plans to offer a broad array of its software catalog as Software-as-a-Service (SaaS) on AWS.

Building on IBM Software being available as-a-Service (aaS) on IBM Cloud, this first-of-its-kind agreement between IBM and AWS will provide clients with quick and easy access to IBM Software that spans automation, data and AI, security and sustainability capabilities, is built on Red Hat OpenShift Service on AWS (ROSA), and runs cloud-native on AWS. The two companies are also committing to a broad range of joint investments to make it easier for clients to consume IBM Software on AWS, including integrated go-to-market activities across sales and marketing, channel incentives, developer enablement and training, and solution development for key verticals and industries such as Oil and Gas, Travel and Transportation, and others.

Today, organizations are looking for industry leading services and solutions that allow them to be nimble, flexible, and continuously scalable. This need has been further compounded as demand grows to run software both on-premises and across hybrid cloud environments so they can be scaled globally with high availability.

Moving forward, organizations will be able to run a broad array of the IBM Software catalog as cloud-native services on AWS so they can get up and running quickly to deliver business value. This includes IBM API Connect, IBM Db2, IBM Observability by Instana APM, IBM Maximo Application Suite, IBM Security ReaQta, IBM Security Trusteer, IBM Security Verify, and IBM Watson Orchestrate, with others to follow later this year.

Clients will be able to procure the IBM SaaS products in AWS Marketplace, and then set up and integrate with AWS services, allowing them to get started with just a few clicks, without deploying, updating or managing any of the infrastructure. IBM SaaS products on AWS are designed to provide high availability and elastic scaling on demand to meet unpredictable throughput needs and will offer a native AWS experience with deep integration of AWS services out of the box and support for API, CloudFormation and Terraform templates to enable automation of end-to-end workflows.

For example, using [IBM Maximo Application Suite as-a-Service](#), a manufacturer will be able to take a flexible, demand-based approach to AI-driven asset management to help them monitor and maintain equipment more efficiently, or predict potential mechanical failures to fix them before they create interruptions. By taking advantage of a scalable consumption model for these applications, they can free up capital for innovation, prototyping, tooling and production – and easily expand their usage over time based on evolving market trends and production demands.

More Articles

[IBM Federal Ecosystem Supports Executive Order Implementation](#)

[IBM Updates Benefits Program for IBMers and Retirees](#)

[IBM Announced as COP27 Technology Partner](#)

[Subscribe to email](#)

Additional Assets

Db2 Four big bets for 2025

Db2 and Db2 Warehouse SaaS

Roadmap evolution including new Enhancements to IBM Cloud and AWS offerings

Available on additional Hyper Scalars in 2025

Db2 infused with Generative AI

We're adding Gen AI capabilities into Db2. Stay tuned.

UX overhaul for management console

Continued investment to improve the user experience for devs and DBAs

Db2

Multiple 12.1.x mod packs planned for 2025

Introducing the AI-powered Database Assistant

Available now via Db2 & Db2 Warehouse SaaS, the new Database Assistant is designed to help you quickly resolve database issues, manage a growing set of databases, and bring new DBAs up to speed faster than ever before.



Db2 Expert

Get answers to your Db2 questions, faster

What is the default page size in Db2?
What is the purpose of buffer pools?
How do I create a range-partitioned table in Db2 11.5?



Monitoring Metrics

Quickly access key Db2 metrics using natural language queries

Provide me the current list of connections
Retrieve CPU usage from September 9th, 2024
Show me storage utilization by tablespace



Simplified Troubleshooting

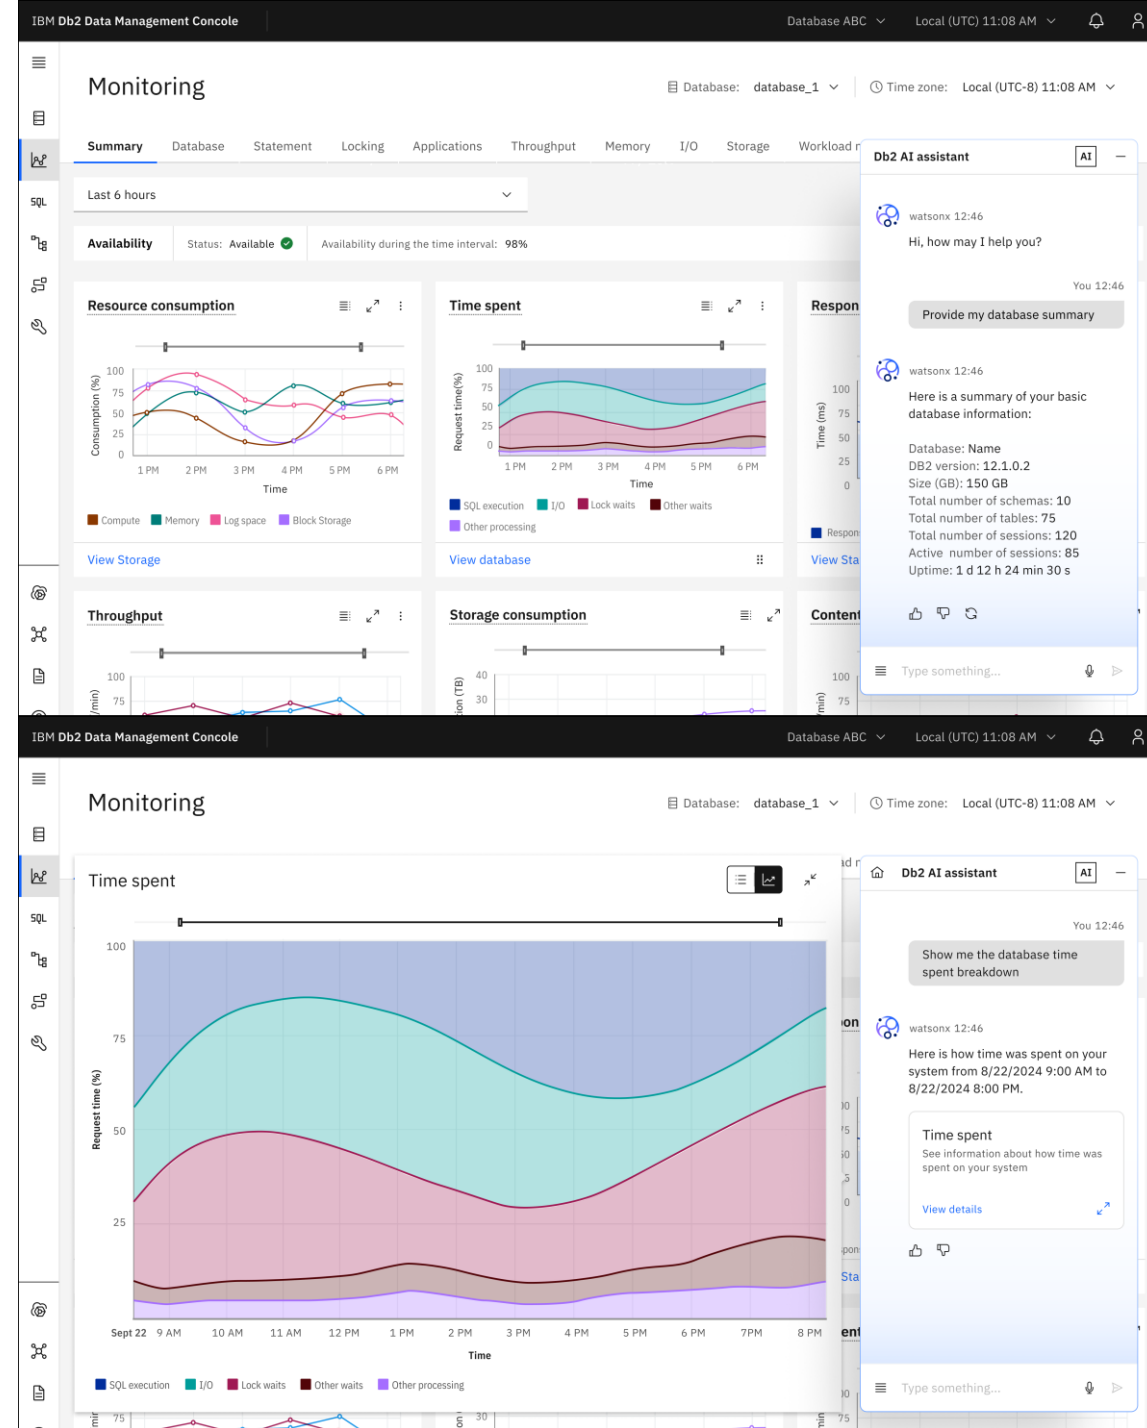
Get recommendations for troubleshooting common Db2 issues

What is sqlcode=-964?
"TABLE1.COL1" is not valid in the context where it is used.. SQLCODE=-206, SQLSTATE=42703, DRIVER=4.32.28



Advanced Analysis (coming)

Identify root cause of performance issues, bottlenecks, deadlocks



Next major release of the Db2 Database Management Console

In 2025, we plan to launch a brand new Db2 management console:

- New** Entry Hub
- New** Monitoring Hub
- New** Data Workbench
- New** Administrator Hub
- New** Conversational Database Assistant

and much more!

The screenshot displays the IBM Db2 Database Management Console interface. At the top, a red banner reads "Roadmap: Dates & content subject to change". The main interface is dark-themed and includes a sidebar with navigation icons for home, SQL, and other tools. The top navigation bar shows "IBM Db2" and user profile icons.

Welcome, user

AI Summary
Database X in OLAP has high activity due to the problem in xyz area. This will affect the following data. Database Y in OLTP has slower activity due to the abc issue. This will later affect areas in 123. Database X in OLAP has high activity due to the problem in xyz area. This will affect the following data. Database Y in OLTP has slower activity due to the abc issue. This will later affect areas in 123.
[View all alerts](#) →

Select database tag
10 database tags
[Manage database tags](#)

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

- All databases 50 databases 1 alert 10 warnings
- Db2 Warehouse 10 databases 14 warnings
- Db2 Database 10 databases 100 alerts 100 warnings
- pureScale 10 databases

All databases
10 databases

Storage: 50 TB
CPU: 20 cores
Memory: 100 TB

Search databases: [New connection](#) +

Name	Tags	Alerts	CPU	Memory	Storage	Total ti...	Q...
database_01	25	2 alerts 2 warnings	2%	12%	18%	1d	100
database_01	25	2 alerts 2 warnings	2%	12%	18%	1d	100
database_01	25	2 alerts 2 warnings	2%	12%	18%	1d	100
database_01	25	2 alerts 2 warnings	2%	12%	18%	1d	100

5 1–5 of 50 items 1 of 10 pages

Notifications
Filter by: All alerts Time frame: Newest

Total activity time
database_01
07/01/2024 08:02:70 PM

CPU time
database_01
07/01/2024 08:02:70 PM

CPU time
database_01
07/01/2024 08:02:70 PM

CPU time
database_01
07/01/2024 08:02:70 PM

CPU time
database_01
07/01/2024 08:02:70 PM

Total activity time
database_01
September 24, 2024 8:02:70 PM

Description
This database has exceeded the threshold for total activity time.

Recommended solution
Throughput has been below baseline by 30% for the past day because of a long running query. Adding an index to optimize this query will improve throughput and overall database performance.

[Open monitoring hub](#) → [Work with Database Assistant](#)

Activity center
10 jobs
[View all jobs](#) →

Job Name	Status	Database
Replication monitoring	In progress	database_01
Cluster health check	In progress	database_02
Content	In progress	database_02



One Data Movement tool for all your mission critical workloads.

"What"

Move any workloads with Db2 being the target

- **IBM Db2 (source/target)**
- **IBM Db2 Warehouse (source/target)**
- **IBM Integrated Analytics System (source)**
- **Power 10 Cloud Rack (target)**
- **Db2 Data files (source/target)**
- **Oracle (source)**
- **Teradata (source)**

"Where"

Support all possible deployment options for Db2

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

"How"

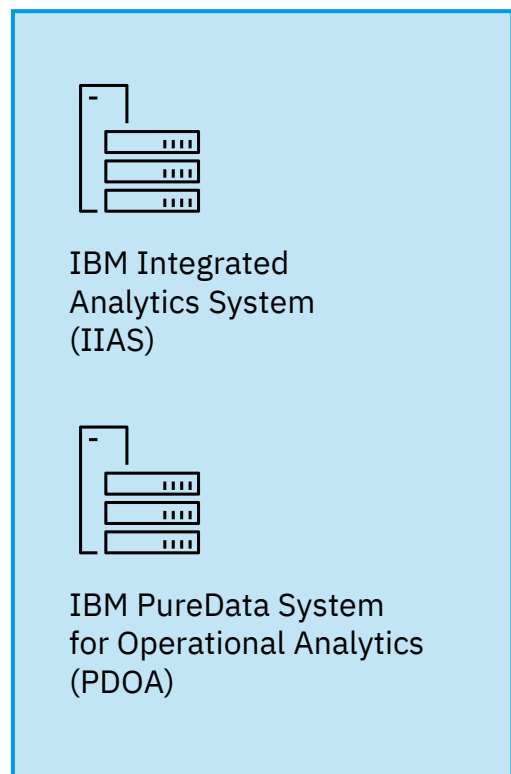
Extensive set of features to accommodate complexities of data movement

- **Schema-level selection of objects**
- **Move massive number of databases or large databases with hundreds of Terabytes or Petabytes of data**
- **Enable High Performance Unload for faster migrations**
- **Support:**
 - **Row/Columnar organization**
 - **Little/Big Endian formats**
 - **All object types, Security labels Privileges, Authorizations**

"When" : Planned for Q1 2025

Db2 Warehouse Offerings

Migration Paths for Existing PDOA/IIAS (Sailfish) Appliances



Migrate to...



1. Db2 Warehouse SaaS

Fully managed, high performance, cloud-elastic data warehouse


- ↳ Runs on both AWS and IBM Cloud
- ↳ Available in US, EU, UK, JP, AP commercial regions



2. Db2 Warehouse Self-managed

Self-managed data warehouse **on cloud provider** managed Kubernetes or OpenShift infrastructure.

- ↳ Runs on Amazon EKS, Amazon ROSA, Azure AKS, Azure ARO
- ↳ Available worldwide



3. Db2 Warehouse Cloud Rack

Self-managed, high performance data warehouse deployable on **on-premises** OpenShift on Power 10 infrastructure

- ↳ Runs on *POWER10* next-gen on-prem offerings

IBM Db2 Warehouse on Power 10 Cloud Rack


Complete Data Warehouse Solution

Simplify and accelerate your data warehouse deployment in days with our Power 10 Based pre-configured solution that includes everything you need to securely analyze your Data on your premises – including Servers, Storage, Switches, and Services!




Simple

Ease into a Db2 Warehouse with a pre-configured, all-in-one design that removes the guesswork




Blazing Fast

Columnar-organized, memory-optimized data warehouse



Scalable & Flexible

Build the data warehouse needed for today and easily scale and adapt as your future needs change



Open for Integration

Support for open formats and integration with watsonx.data

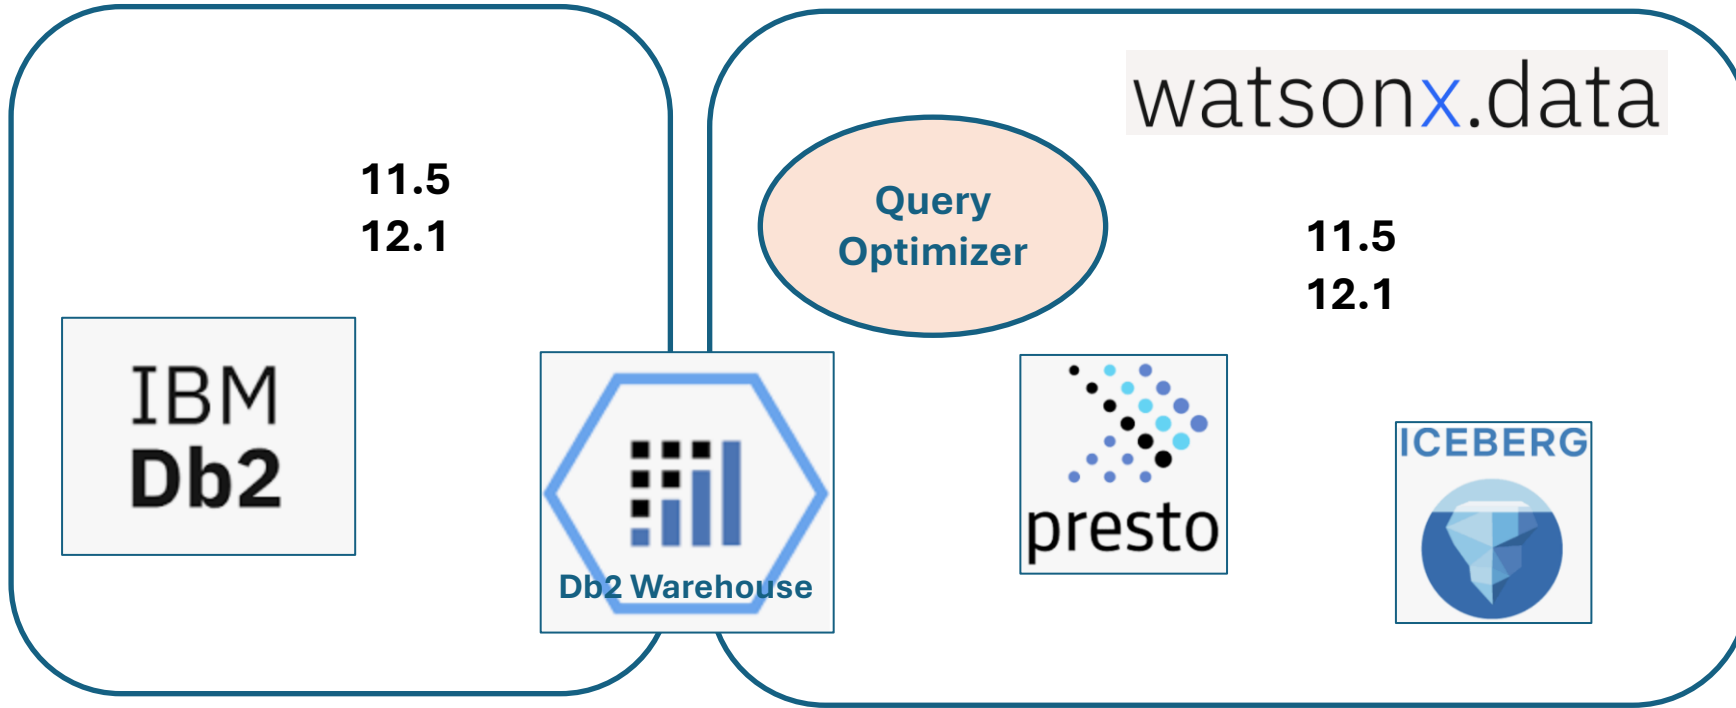


Durable & Reliable

Industry-leading continuous availability and disaster recovery

Db2 Universal Container


Modernize your Db2 workloads anywhere !



Db2U is the modernization delivery vehicle for:

- Db2 and Db2 Warehouse on Cp4D/Open Shift/K8s
- Db2 Warehouse on Cloud Gen 3 (AWS) and IBM Cloud (VPC Gen 2)
- Hyper-scalers (AWS, Azure and IBM Cloud)
- Next-gen Open Architecture for Appliance Modernization (IBM P10 Cloud Rack)
- Integrating with watsonx.data ecosystem including Lakehouse (DATA LAKE table) support

Platforms



Db2 12.1

Db2 12 brings significant enhancements to Db2 pureScale, name space separation, generative AI-powered insights, a new AI optimizer and hundreds of other enhancements.

Db2 pureScale and HA integration improvements

Complete replacement of TSA with Pacemaker technology for cluster management in Linux, leading to significantly faster failure recovery times and better user experience

AI-powered query optimizer

Allows Db2 to continuously learn from customer's queries and achieve up to 3x query performance improvement over prior version

Name space separation with TENANT construct

Create a logical separation between one or more database schemas, easily isolating differing sets of tables from each other

Db2 infused with Generative AI

We're adding Gen AI capabilities to Db2. DB Assistant and more coming in 2025.

- **Improvements to backup performance by initiating multiple threads to process a single table space**
- **Mac M1/M2 driver support** for developers on macOS using Apple Silicon chip
- **HADR upgrade leverages RoS for enhanced availability**
- **Db2 pureScale HADR support for enterprise-grade end-to-end SSL encryption**
- **Online index reorg for Db2 pureScale** allowing index reorg while table remains online/available (mod 1)
- **ADMIN_MOVE_TABLE** performance and availability enhancements
- **Security enhancements** with AUDIT exceptions, Trusted Context and data masking
- Continuing investment in **cloud object storage performance**
- **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
- **Federation enhancements** with support for Snowflake, Oracle 23c and performance improvements



Db2 12.1 Supported OS and Platform

Distro	Version	Db2 Server	Db2 Client/DSDRIVER	HA	pureScale
RHEL	9.4	x86-64 System z Power 9,10 LE	x86-64, x86-32 System z (64bit only) Power 9,10 LE (64bit only)	x86-64 System z Power 9,10 LE (Pacemaker)	x86-64 ★ System z Power 9,10 LE (Pacemaker)
SUSE	15SP6	x86-64 System z Power 9,10 LE	x86-64, x86-32 System z Power 9,10 LE	x86-64 System z Power 9,10 LE (Pacemaker)	x86-64 ★ (Pacemaker)
Ubuntu	22.04 LTS	x86-64 System z Power 9,10 LE	x86-64, x86-32 System z Power 9,10 LE	N/A	N/A
AIX	7.3TL2	64bit Power 9,10 LE	64bit Power 9,10 LE	64bit Power 9,10 LE (TSA)	64bit ★ Power 9,10 LE (TSA)
Windows Desktop	11	x86-64	X86-64, x86-32	N/A	N/A
Windows Server	2022	x86-64	X86-64, x86-32	N/A	N/A
Mac	Sonoma ★ M1/M2/M3 Arm chip	N/Z	64bit DSDRIVER only	N/A	N/A
UBI Open shift	9 ★	x86-64 System z Power LE	N/A	x86-64 System z Power LE	N/A

SLES pureScale delayed to 12.1.1 due to dependency delays

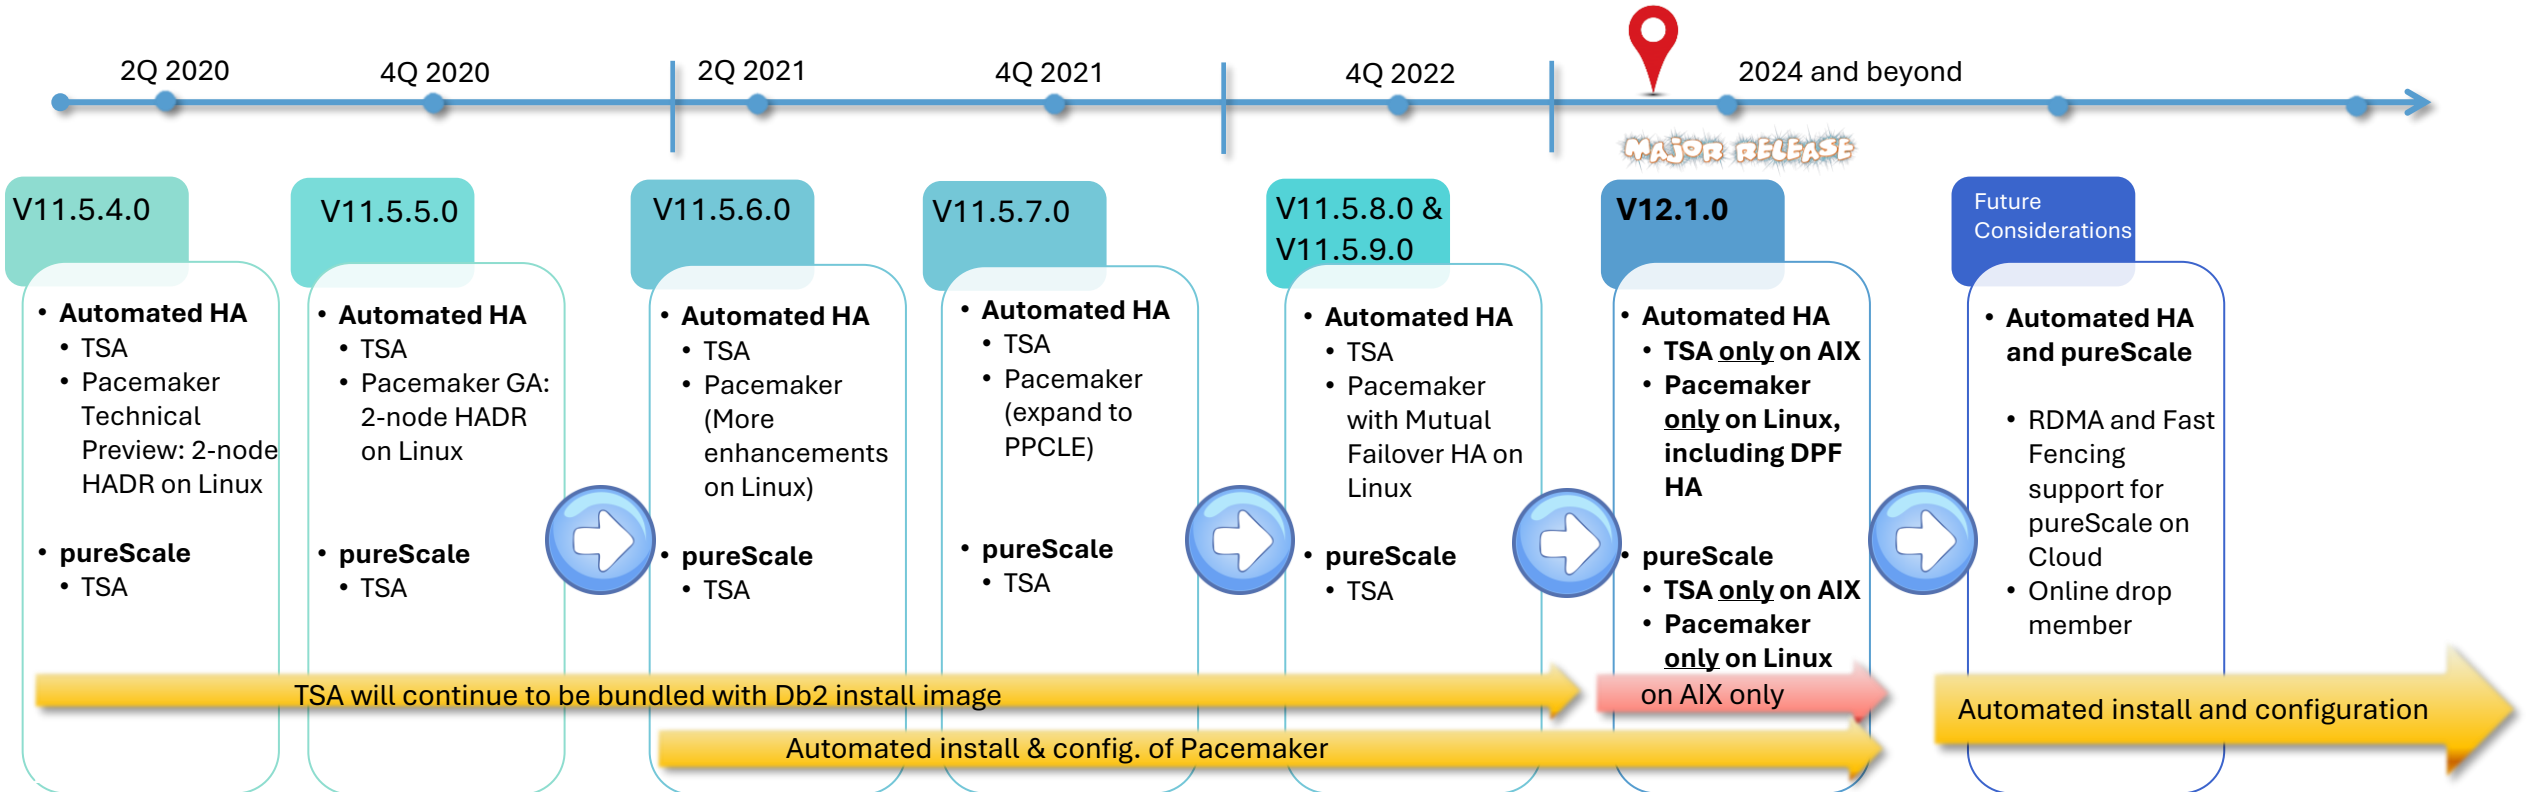
Pacemaker Integration with Db2

**A
D
V
A
N
T
A
G
E
S**

- 18+ years in industry as HA resource cluster manager
- Included by RHEL and SuSE as paid add-on HA package
- Open source: inspect, collaborate and innovate
- Align with IBM Open-Source Strategy

**B
E
N
E
F
I
T
S**


<p>Cloud Ready</p>	<p>PERFORMANCE Faster Failover Recovery</p>	<p>Lower \$\$\$ Cluster Software</p>
<p>First contribution to community in V11.5.6.0!!!!</p> <p>STACK MODERNIZATION open source</p>	<p>One Team Just Db2 Technical Support</p>	<p>Simpler Architecture</p> <p>Simplified PD</p>




Note: Roadmap subjected to change

Pacemaker Integration with Db2


Single DB Partition (EE) with automated HADR



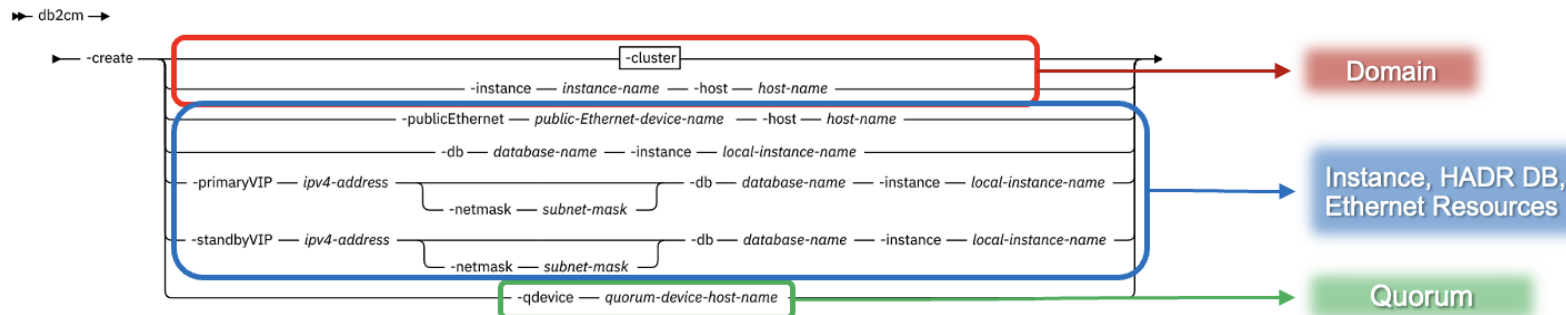


Mutual Failover (a.k.a. Active/Passive) with automated HA with shared storage



Database Partitioning Feature (DPF) with automated HA (same site)

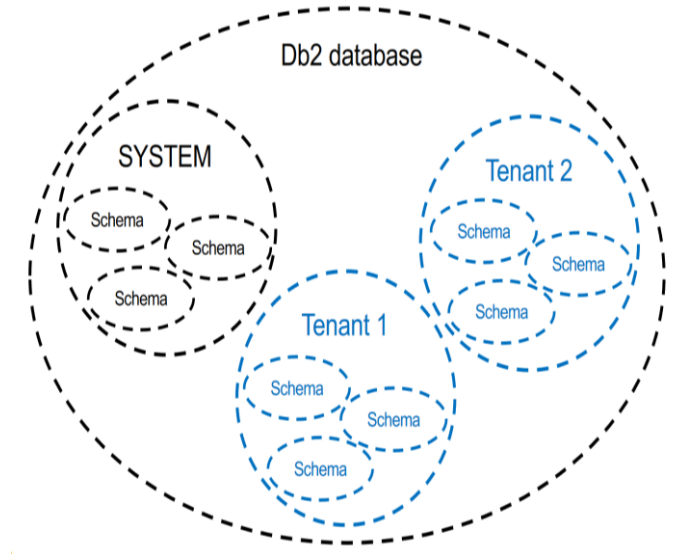


pureScale Online 24x7x365 with automatic failover



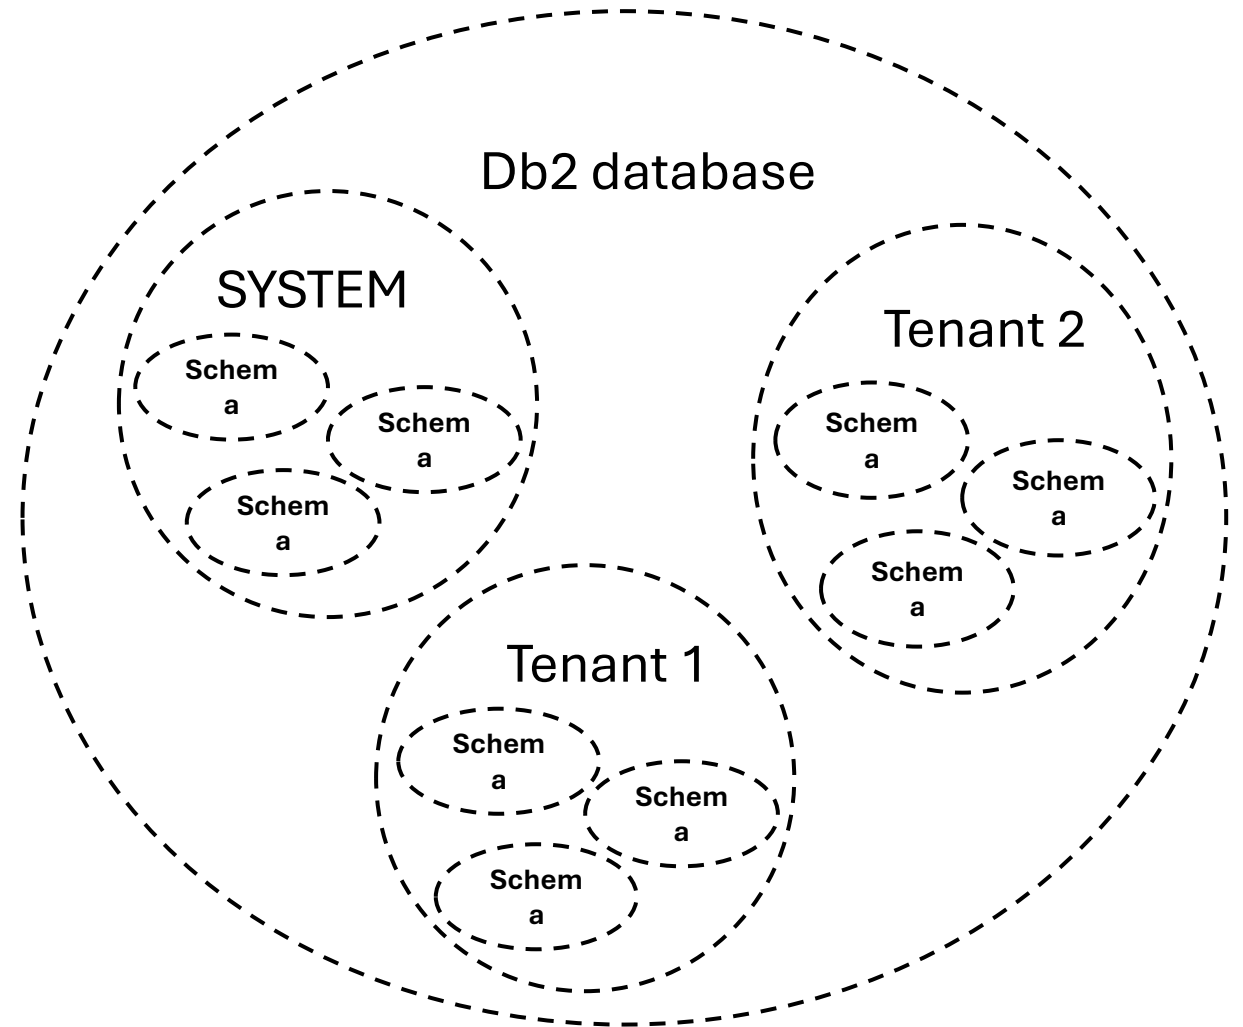
Db2 Namespace Separation (1|2)

- What is Namespace Separation (Tenant)?
 - A single physical database supporting multiple “Tenants”
 - All tenants share the same infrastructure and resources
 - Each tenant has a private, isolated perspective for their own objects
- What problem does having Tenants solve?
 - ✓ Netezza/PostgreSQL compatibility
 - ✓ Cost-savings through consolidation
 - ✓ Reduction of fixed overhead costs associated with individual databases
 - ✓ Centralization/simplification of database operations
 - ✓ Sharing unique environments without collisions
 - ✓ Significant cost-savings for our customers with many small development systems



Db2 Namespace Separation (2|2)

- The default SYSTEM tenant
 - The initial set of catalogs established when the database is created
 - This is referred to as the SYSTEM tenant
 - Cannot be removed
 - Contains catalog information for shared resources and Db2 defined objects
- User-defined tenants
 - A DBADM can create a tenant to set up an independent catalog namespace within a Db2 database using CREATE TENANT statement
- All connections made to a database are initially associated with the SYSTEM tenant
 - A SET TENANT statement must be issued to associate a connection with a user-defined tenant



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.



Cardinality estimation for a subset of local predicates using a learned model



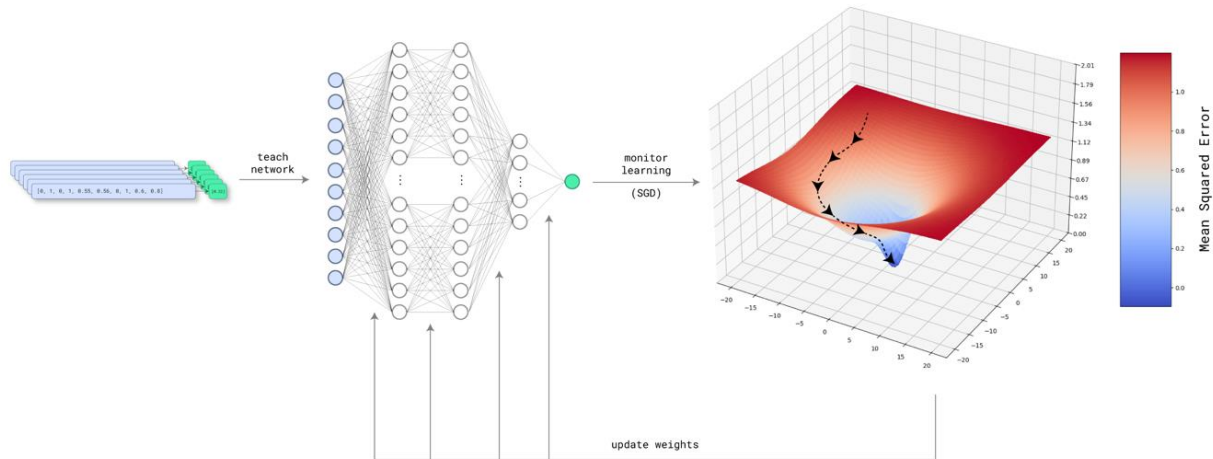
Model policies to filter which tables can have models discovered on them



Explain support to show model information in the execution plan



DDL to disable/enable, revert, and drop models



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



Db2 AI Query Optimizer

CHD.PAY_0 IN (0,1,2) AND
 CHD.PAY_2 BETWEEN 0 AND 2 AND
 CHD.BILL_AMT1 BETWEEN 150 AND 746814
 AND
 ...

Algorithm	Cost	Rows	Tables
Traditional Algorithms	4838.66 (4.5729e+06)	4	HSJOIN
Traditional Algorithms	7.9354 (15032)	8	HSJOIN
Traditional Algorithms	10.102 (15032)	12	TBSCAN
Traditional Algorithms	7001.44 (50245)	106	TBSCAN, CO-TABLE: DEMO SENTIMENT_SCORE_DATA
AI Model	37787.2 (4.5729e+06)	5	HSJOIN
AI Model	61.971 (7516)	7	HSJOIN
AI Model	54677.3 (50245)	8	TBSCAN, CO-TABLE: DEMO PURCHASE_HISTORY
AI Model	79.5008 (82)	9	HSJOIN, CO-TABLE: DEMO PURCHASE_HISTORY
AI Model	114.982	10	TBSCAN, CO-TABLE: DEMO CREDIT_HISTORY_DATA
AI Model	114.982	11	TBSCAN, CO-TABLE: DEMO SENTIMENT_SCORE_DATA
AI Model	114.982	16	TBSCAN, CO-TABLE: DEMO POLICE_DATA



MODEL AUTO CGS INDEPENDENCE ASSUMPTION

Security Enhancements in Db2 12.1 (1|2)

- Trusted Context / Connection
 - Audit Exceptions for trusted applications to limit auditing
 - Define a Trust procedure to implement custom logic to identify Trusted Ctx
 - Generate an AppToken to identify a custom application
 - Role can be used for package execution/static SQL ([Aha! 407](#))
- Improve update/upgrade with RCAC dependency on system objects (catalog views, functions etc.) ([Aha! 1004](#))
- JWT authentication exposed as Global Variable for use with RCAC (JWT available since 11.5.4)
- Behavior Changes
 - DBADM extended to support tablespace and buffer pool creation
 - DBADM can create automatic storage tablespaces defined on SYSADM created storage groups
 - DBADM can create buffer pools (instance_memory still set by SYSADM)
 - DBADM no longer granted DATAACCESS by default.
 - PUBLIC no longer granted by default on database creation
 - CONNECT, IMPLICIT_SCHEMA, CREATETAB, BINDADD, USE on USERSPACE1, CREATEIN on SQLJ and NULLID schemas
 - No Change existing databases ie. No revokes issued on upgrade.
 - Hostname Validation ON By Default

Security Enhancements in Db2 12.1 (2|2)

- TLS / SSL Enhancements
 - Change to KMIP SSL label now an online operation
 - TLS 1.3 – on by default at server
 - TLS + SERVER auth can be used to connect in place of SERVER_ENCRYPT authentication (SERVER_ENCRYPT deprecated)
- Multiple behavior changes related to FIPS, to leverage major release boundary.
 - E.g. SERVER_ENCRYPT is blocked in STRICT_FIPS mode
- Discontinuations:
 - TLS 1.0 and 1.1 removed from Db2 server
 - Also removing support for SHA-1 and 3DES cipher suites
 - 3DES with native encryption
 - Existing database & backups continue to decrypt.
 - New databases & backups needs to be AES (Reading from old 3DES backups still ok)
 - DATA_ENCRYPT authentication
 - Replaced by TLS/SSL
 - CLIENT Authentication
 - FED_NOAUTH DBM CFG
 - Change password plugin
 - Default OS plugin provides equivalency. Legacy from earlier implementations

Columnar Improvements in 12.1

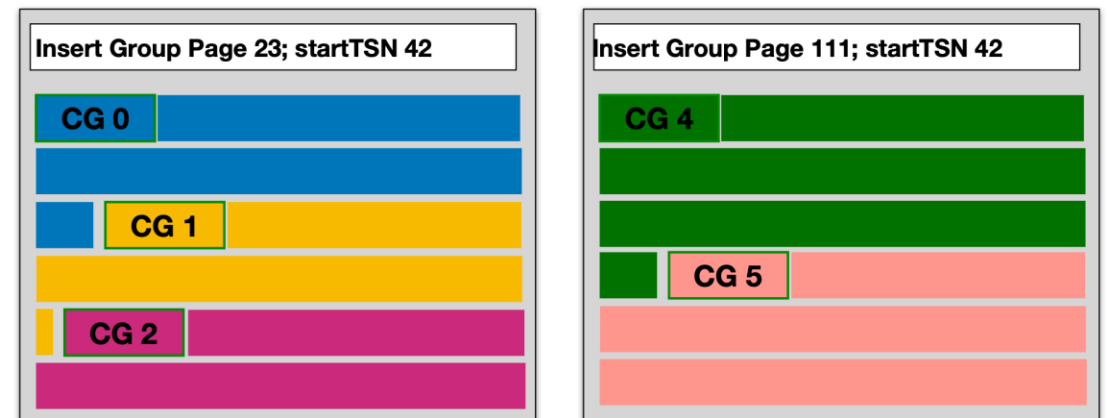
- Enhanced Insert and Compression enabled by default
 - Page based string compression: very effective for compressing high cardinality hex, date, timestamp, numeric data stored in string data types
 - Trickle Insert performance improvements
 - Deferred Synopsis creation for small tables
- Schema Evolution: DROP and RENAME columns now supported
 - ADMIN_MOVE_TABLE – online columnar table moves
 - V12 will see schema evolution evolve throughout the mod packs
- Compact Varchar Extended to Hash Join
 - Individual queries saw performance improvements of 5x – 60x!!
 - Overall memory requirements reduced by 75% - 80% (4x -5x “less”)

Columnar Improvements in 12.1

Trickle Feed Insert Enhancements

- Used only when small number of rows are being inserted (aka data trickling in).
- Inserted rows are split to one or more “insert groups” – still columnar format just inserting more columns per page.
- Number of insert groups depends on types of columns, average length, etc. But generally, will be much less than total number of columns.
- Data going into these insert groups are not compressed.

Insert Groups




- Flexible assignments of column groups to insert groups.
 - Fixed-length vs. variable-length, large varchar columns.
 - Co-existence with other approaches, e.g., text compression.

Columnar Improvements in 12.1

Trickle Feed Insert Enhancements

- The insert group pages are almost always temporary -- a window of the most recent 'trickle' inserts. Exception is small tables.
- Insert group pages/rows are automatically moved (aka split) to the single column per page format.
- The split is done synchronously. Triggered as soon as it's predicted that full column group pages can be created.

Split Insert Group



- **Synchronous** within a transaction (same approach as used for page compression).
- Tries to split when it thinks that a column group page can be filled.
 - Very **few pages** are **in Insert Group format**.
- Each column group is processed independently.

Columnar Improvements in 12.1

Compact Varchar (CVC): Recap of the journey

- Why?
 - Improved memory stability and performance of the columnar engine
 - Specifically for workloads involving queries with VARCHAR columns
 - Dramatically reduce Out of Memory (OOM) / -955C errors
 - Improve individual query and overall workload performance
 - Reduce need to modify schema to better size varchar columns
- Phased approach:
 - Vector and Work Units ✓
 - Group by and Join ✓
 - Aggregation and OLAP ✓
 - M to N Joins: new

Columnar Improvements in 12.1

Compact Varchar Hash Join

- CDE Hash Join improvements:
 - Used for Equality Joins in CDE
 - Reduce memory requirement for large “Payloads”
 - i.e. Inner/Dimension tables columns
 - Improve performance for joins with large “Payloads”
 - HSJoin converted to use new “Compact Block Store” (CBS)
 - Variable length inners and wide inners are stored in the CBS
- CVC Extended to cover all variable length data types
 - VARGRAPHIC, VARBINARY, LOB descriptors

Columnar Improvements in 12.1

Compact Varchar Hash Join

- Evaluation:
 - Variable length “Payloads”:
 - “Payloads” are smaller because they are stored compacted
 - Performance improved overall by 2.4x - 2.5x across a wide range of Sort Heap sizes (600K, 300K, 150K, 75K pages)
 - Individual queries saw performance improvements of 5x – 60x!!
 - Overall memory requirements reduced by 75% - 80% (4x -5x “less”)
 - Many queries no longer spill, leading to the huge improvements
 - Fixed Length Payloads
 - “Payloads” are the same length, but handled as a “blob” by Hash Join
 - Performance improved overall by 34%– 59% across a wide range of Sort Heap sizes (600K, 300K, 150K, 75K pages)
 - Individual queries saw performance improvements of 3.7x – 4.5x!!
 - Overall memory requirements reduced by 20% - 50% (26% - 2x “less”)

Database Upgrade Availability

HADR Read Access on standby during upgrade


- HADR now allows read access on standby during primary upgrade
 - New upgrade procedure: upgrade primary database first while keeping standby database at v11.5 to provide read only access
 - After primary database has completed the upgrade, move applications back to the primary then upgrade the standby database
 - Trade-off: read access during upgrade of primary vs no HA protection until standby completes upgrade and reaches PEER state.
 - If multiple standbys, the following order to avoids gap
 - Upgrade principle standby, then primary, while keeping aux standby for read-access
 - Upgrade aux standby after upgrade of primary is complete

Backup Performance

Intra-tablespace parallelism (1|2)

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




Backup Performance

Intra-tablespace parallelism (2/2)

Single node with 2 FC Cards:

- 3.8X faster in extreme skewed case, best scenario.

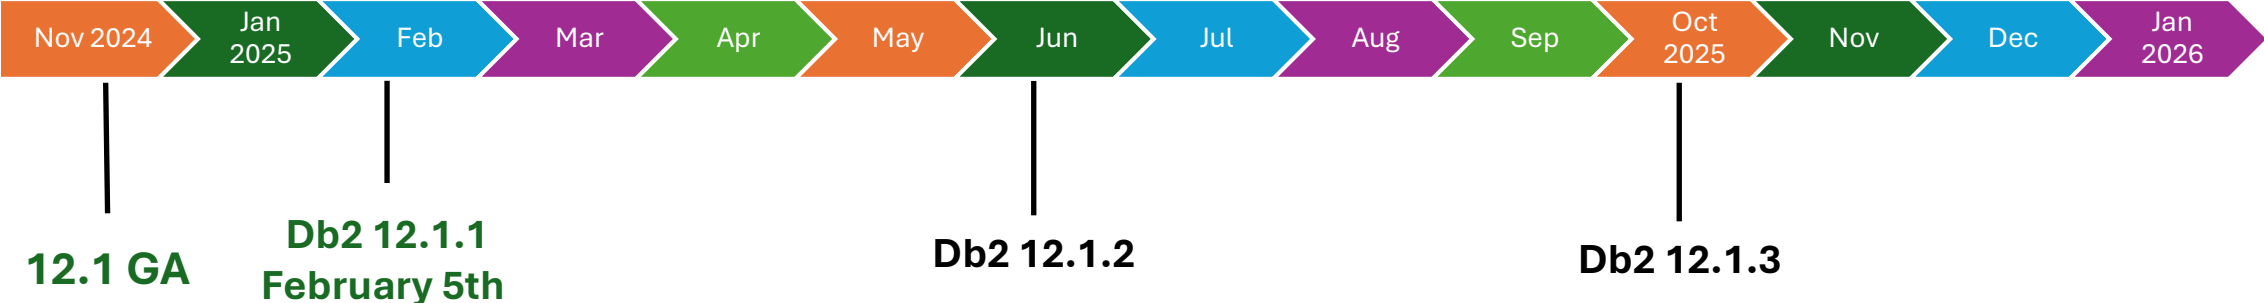


Things to be aware of

- Linux pacemaker will be the only integrated cluster manager
 - including pureScale
 - AIX remains with TSA
- ICU only latest version will be carried forward
 - mod packs may update ICU versions – ICU 74 ships with GA

Db2 12.1 mod pack our current thinking

But subject to change



12.1.X Candidates

12.1.1

SVM uplift:

- SLES 15 SP6 – pureScale support
- NVMe reservation

Online Index Reorg for pureScale

AI Optimizer Explain Additions

Columnar:

- FFNR performance improvements
- Explain Enhancements

TBSPACEADM

- Bufferpool & Tablespace authority

12.1.2

Vector and Similarity Search

pureScale: AWS Elastic Fabric Adapter (EFA) support

Columnar schema Evolution: Decimal Precision

World Writeable File Permissions Removed (optional)

Data Masking at read

Remote Storage: Azure Support

12.1.2+

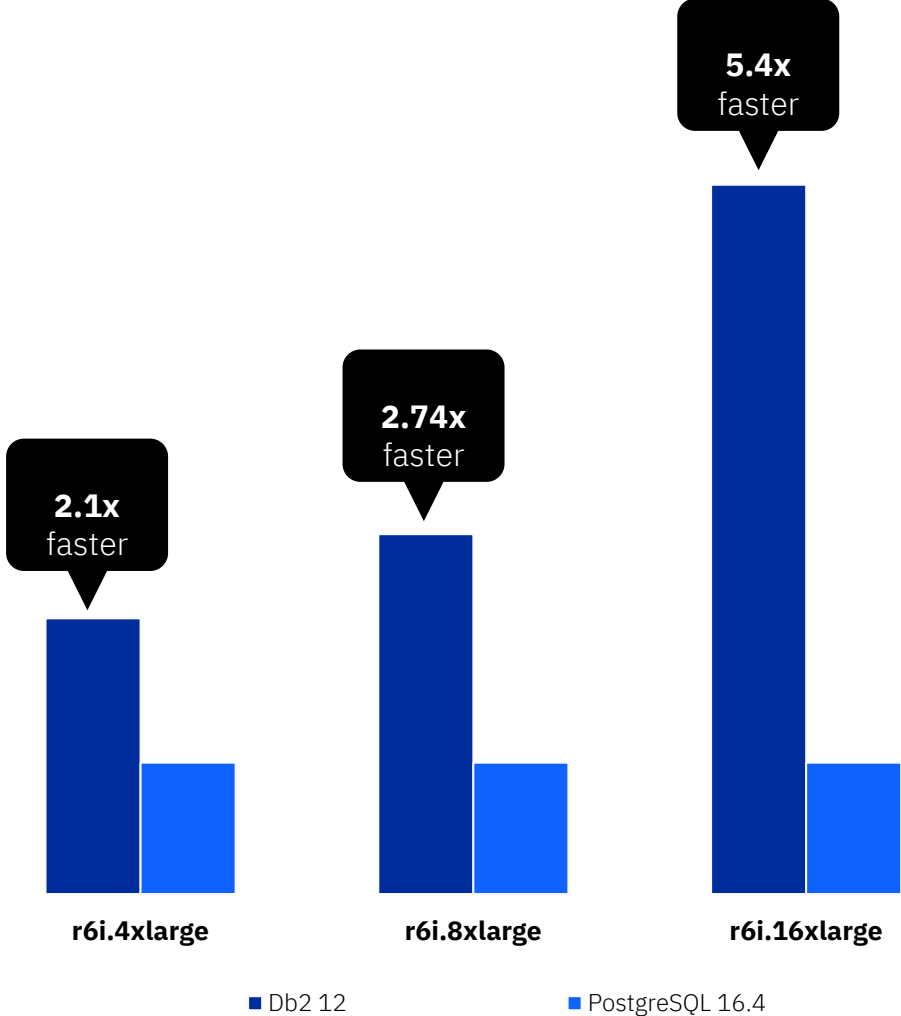
Reorg enhancements

- Online Reorg with Global Index
- pureScale: Online topology Changes
- Remote Tablespace Support
- PPCLE support
 - Externalized to non-container
 - General performance enhancements

One Last Thing

Db2 12 vs PostgreSQL 16.4

5.4x
faster
than
PostgreSQL



Unofficial TPC-E benchmark for OLTP workloads run on AWS EC2 infrastructure


TPC-E Configuration

- 40k, 80k, 160k customers as a function of db size
- 40, 50, 80 connections
- 60 trading days of data

AWS EC2 Instances

- r6i.4xlarge: 16 vCPUs / 128GB mem
- r6i.8xlarge: 32 vCPUs / 256GB mem
- r6i.16xlarge: 64 vCPUs / 512GB mem

Thank You



- Mike Springgay
IBM, Db2 Chief Architect
- springga@ca.ibm.com

IBM