

What's new with Db2?

IBM
Db2

Midwestern Db2 User Group
December 2021

Tanmay Sinha, Lead Product Manager, Db2
@tanmaysinha 



Agenda

Db2 Overview

Feature Highlights


Platform integrations

Roadmap & Next Steps

Q&A

Overview

Db2 Distributed 11.5 – Journey so far...



*Quarterly updates to Db2 Cloud Services
Db2 on Cloud and Warehouse on Cloud. Db2WoC is available today on AWS as a fully-managed service integrated with Cloud Pak for Data*

One converged database, *your choice of deployment.*

Transactional

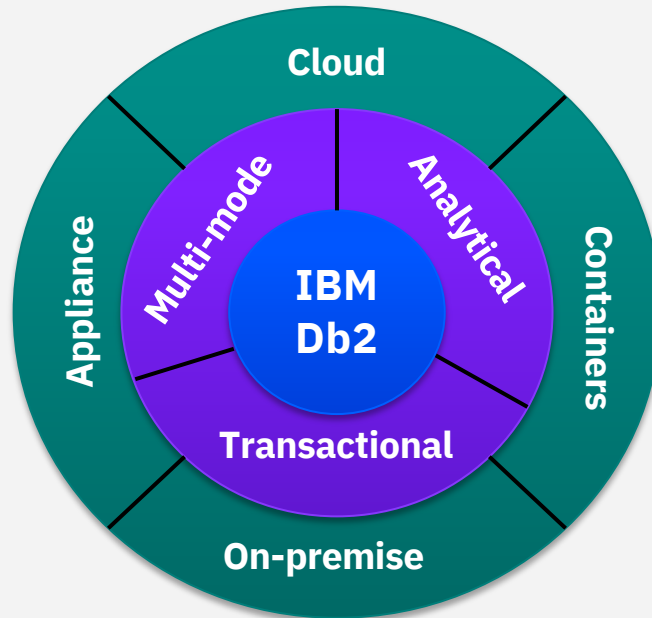
Real-time execution of large volumes of database transactions.

Analytical

Multi-dimensional analysis on large volumes of data from a centralized data warehouse.

Multi-mode

JSON docs, streaming data, graph queries, spatial analytics, and ML workloads.



Cloud-First

Same Db2 available on the platform of your choice. On-premise, on-containers, on-cloud.








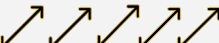
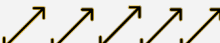



Perfect Compatibility

Full application transparency. Migration tooling to automate and accelerate upgrades.

Design Driven

Persona based experiences for database administrators, application developers, and now data scientists. Consistent experience across all deployments.

One Db2 Engine – Multiple deployments

	On-premise	Containers	Fully-managed
Cost of Administration			
Time to Deploy			
Scalability on Demand			
Configurability			

Feature Highlights

Always on – No outages with active-active configuration

- _Based on **IBM Q Replication** log-capture/transaction replay technology
- _Integrated with the **IBM Analytics Systems appliance** and **available with Db2 LUW** for On-Premise deployment

- _**Continuous availability** of warehousing applications on large analytics systems, using Active-Active configurations
- _Supports for **dissimilar databases** that may reside in different data centers

- _Setup replication between IBM Cloud and AWS with Db2 Warehouse on Cloud
- _Fully managed service available on **IBM Cloud** and **AWS**



Geo-replicated disaster recovery

- _On-demand Disaster Recovery node in another region
- _In the rare case of region outage **shift their workload to a different region**

- _In addition to failover functionality, the DR node also allows “**read-on-standby**” and improves read-availability.
- _**Automate failover** using APIs or failover to the DR Node and fail back to the Primary node with a few clicks.

- _**Add High Availability (HA)** plan combined with offsite HADR geo-replicated node is the ideal option for **maximum uptime**



Use Db2 Graph to analyze your existing data without moving to another store

- _Enterprise grade **in-database Graph solution** – *now included with your Db2 licenses*

- _Queries data and relationships with existing relational tables, supporting graph analytics and traditional SQL on the same data

- _**No change in Db2 data**, structure or performance. Existing applications are uninterrupted. Db2 is untouched

- _Data is fetched by Db2 Graph at time of execution for data consistency and **real-time updates**

- _Based on **open-source Apache TinkerPop** graph framework using Gremlin as graph query language

- _New optimization strategies that improve performance when navigating across edges


Platform Integration

Cloud Pak for Data – Integrated platform for AI journey

_Delivers a single transformative platform - Modular operating system of interoperable services connecting all data with **simplicity**, faster **adaptability**, **reduced** reliance on **specialization**

_Unlike the traditional mix of many isolated tools, using a platform grants: **unified experience**, integrated **automation of manual tasks**, enforced **governance and security**.

_Cloud Pak for Data is the new way to solve business problems as a hybrid-cloud Data and AI solution that **runs anywhere** and connects **data everywhere**.



IBM Data Fabric with Cloud Pak for Data

Simply put...It provides simplicity.

8X faster, 1/2 cost

Run [Analytics 8x faster](#) than Snowflake at half the cost, in a secure environment that doesn't require copying or moving data.

Data Sources, Types and Domains



Business Apps



Custom Apps



Point of Sale



User Behavior

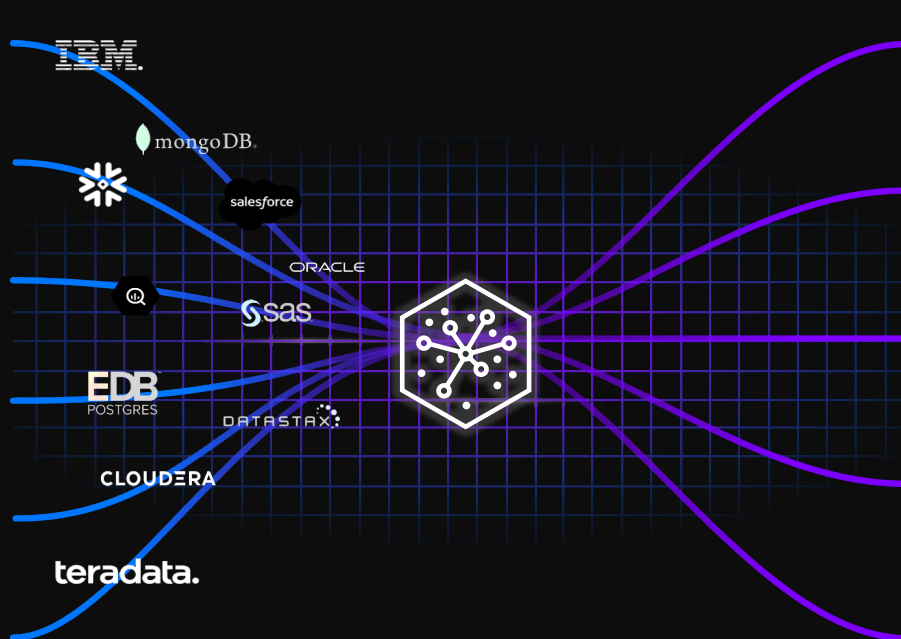


IoT



Devices

Any Data, Any Cloud, Anywhere



Data and AI Outcomes



Customer Centricity



Operational Agility



Total Quality Management



Continuous Improvement



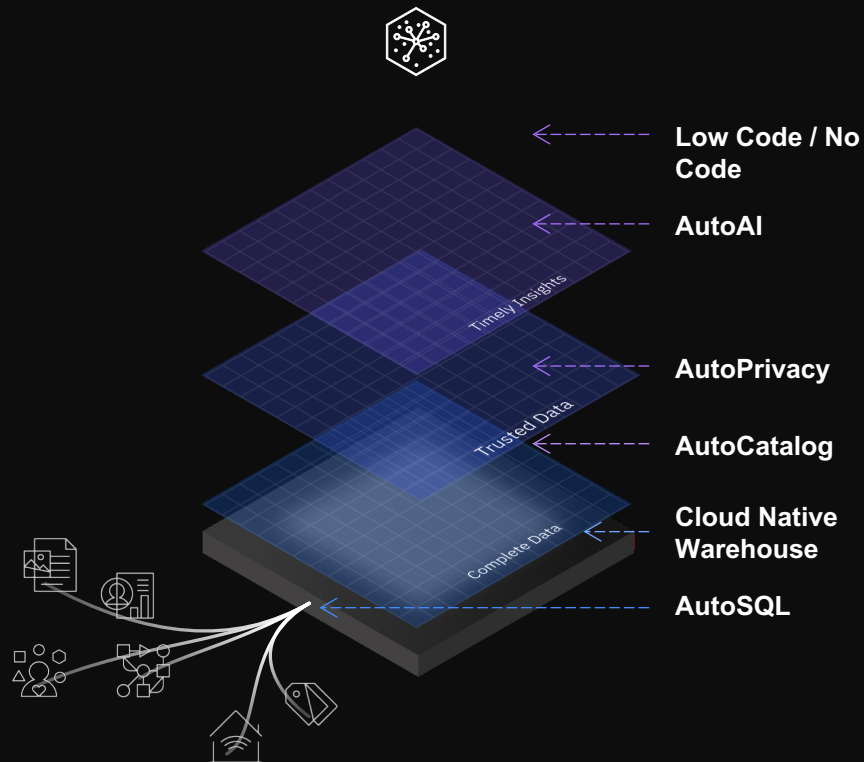
Support Critical Services

Execute with efficiency with the Data Fabric

When you can act on data holistically, you unlock its efficiencies.

53% faster

Data access [with AutoSQL](#) compared to industry benchmarks

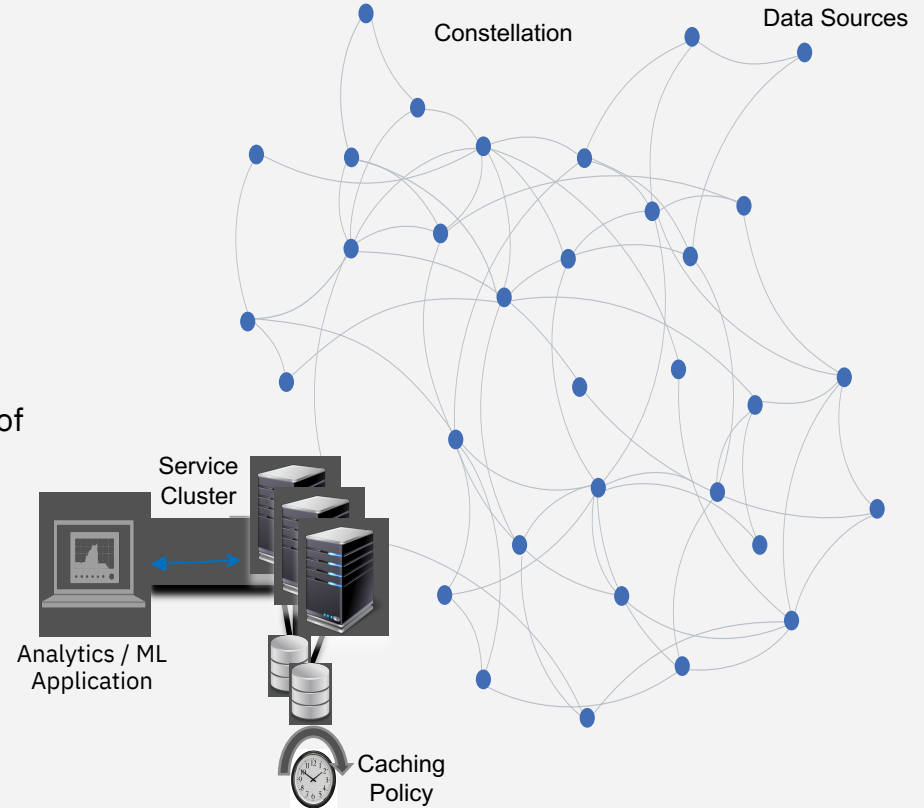


Modern logical data warehouse with Data Virtualization

- _ **Remove data silos** and manage data growth, enable data governance, security and self-service capabilities
- _ **Reduce time-to-value** for new business use-cases and initiatives

- _ Based on **common Db2 SQL** engine and supports rich set of SQL dialects, application portability
- _ Uses a parallel processing mesh providing **execution performance** and **greater scalability**

- _ Retain the use of **existing applications and tools** within the business
- _ Deeply **integrated with Cloud Pak for Data** with access to Watson Knowledge Catalog, Cognos and Watson Studio



Built-in governance catalog across all data assets

_Governance of data and AI lifecycle to **enable trust** across all data in the catalog with **policies automatically enforced**

_Business glossary used to describe information assets across the platform with **reduced cost of custom integrations** with disparate tools

_ **Gain insights fast** with one experience for structured/unstructured data discovery

_ **Automated delta discovery** and bulk term assignment for similar columns.

_Trace data **provenance and transformations**

_Connect data through **lineage charts** with ability to filter the data to show technical vs business vs. custom pathways from source to target.

The logo consists of a large blue circle with a thick blue border. Inside the circle, the text "Watson Knowledge Catalog" is written in blue, and below it, "Quality | Governance | Consumption" is written in a smaller, black font.

Watson Knowledge Catalog

Quality | Governance | Consumption

Top 3 reasons to deploy Db2 with Cloud Pak for Data

Operational Efficiency

Automates and **speeds up administrative functions** such as instantiation, upgrade, failure recovery to greatly reduce the cost.

With Db2 Management Console (DMC), administrator can **monitor their entire Db2 estate** – on premise and on cloud.

Portability

Inherits the flexibility of the underlying OpenShift platform to **deploy on the cloud infrastructure of choice** including AWS, Azure and GCP.

Seamless Integration

Natively integrated with several base services including Watson Knowledge Catalog, Cognos Analytics and Watson Studio.


Expand into new use cases with simple clicks including business analytics, **data science and AI**

Performance results of Db2 Operator with Cloud Pak for Data

Traditional Database deployment

Weeks of Planning and Configuration

Manual install of multiple components including unified console, LDAP, Db2 Rest, Graph API, Data Replication
Additional configuration to get them to work with each Other



Db2 on Cloud Pak for Data

Auto-provisioning and deployment in **< 3 mins**

Fully integrated plug-and-play environment
Upgrade in minutes to latest and greatest without interruptions
Easy to configure automated failover and high availability

Monitor entire Db2 estate with the integrated management console

Design-driven **unified experience across Db2** to monitor hundreds of databases from a single screen
Uses machine learning (ML) algorithms to analyze some issues and propose auto-fixes.

Extensible & composable services architecture with documented **RESTful services**
Explore and manage database objects, such as tables, schemas, or stored procedures with the ability to **write and run SQL** in an integrated SQL editor

No additional cost or licensing restrictions; the console is included with all editions of Db2 and the Db2 family
Available on-premise, on cloud and on Cloud Pak for Data

IBM Db2 Data Management Console				
Databases				
Total: 17 Available: 8 Disconnected: 7 Not monitored: 2 Alerting configuration				
Filter by: Alerts Tags Metrics Clear filters				
Connection name	Responsiveness		Throughput	
	Response time	Statements	Rows read	
yu-wow	93 ms	63		
rui04091	16 ms	6		
rui0320	12 ms	6		
rui0331	11 ms	6		
rui04093	8 ms	6		
sqlnew1154	5 ms	6		
hlconnection	--	--		
hldbnew	--	--		

Roadmap & Next Steps

Db2 11.5.7 – December 8th, 2021



Availability & Recoverability

- Logical Backup support for TSM (incl: IIAS)
- Add know to determine restore success or failure if tablespace error detected (RFE)
- Pacemaker for PPC LE (RHEL 8.2+ and SLES 15SP3+)
- Pacemaker: 2-node HADR support on PPCLE
- Pacemaker: Automated HADR Support
- HADR Database Hang Monitor Enablement for AIX



Core Engine

- RENAME TABLE with INDEX ON EXPRESSION
- Cloud Object Store – Phase 1 – Enable COS storage S3 vendor SDK
- Control File Recovery - Avoid single point of failure for Db2's critical files (RFE)
- ADMIN_MOVE_TABLE for INDEX ON EXPRESSION
- Upgrade to Java 8.0.6.35
- Db2 Advanced AU (SAP)
- AIX 7.3 Support

Db2 11.5.7 – December 8th, 2021



pureScale

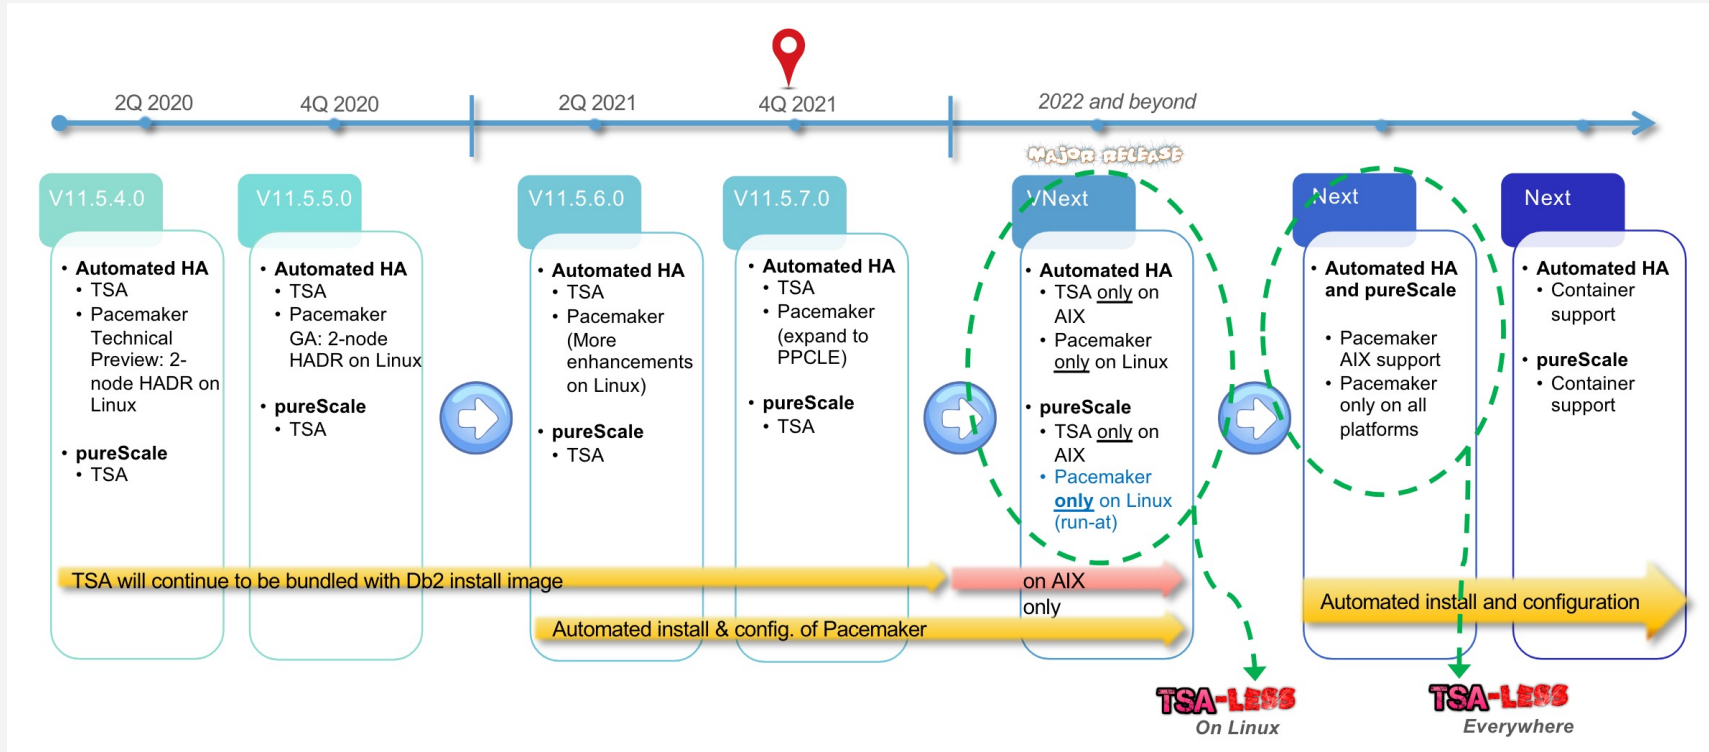
- pureScale Support for zLinux
- Spectrum Scale 5.0.5.8 Support
- TSA 4.1.0.7
- Pacemaker HADR Enhancements
- HADR Enhancements



Appliance

- TSM Integration
- Spectrum Scale 5.0.5.8 Support
- Support for Python 3 for scripts (CC)

High-availability Roadmap




Watson Query – Available 4Q'22

Consolidate & Simplify



Simplified Onboarding



Single deployment experience



Consolidated Resource Management

Name	Type	Created by	Created on	Status
document1	doc	admin	2013-10-20T15:00:00	Active
document2	signed	admin	January 10, 2014	Active
document3	image	admin	January 20, 2014	Active
document4	worddocument	admin	January 20, 2014	Active

A unified resource
provision, monitor, and
manage experience



Unified Query Experience

[illegible]

Single SQL query interface
and experience



Consolidated Administration

The screenshot shows the IBM Cloud Pak for Data console. At the top, there's a navigation bar with 'My workspace' and 'workspace-1'. Below that, there's a breadcrumb trail: 'Data Science' > 'Data Science' > 'Data Science'. The main content area is divided into two sections. The top section, titled 'Overview', contains a table with columns 'Data Source', 'Data Science', 'Data Science', and 'Data Science'. The bottom section, titled 'Data Science', contains a table with columns 'Data Science', 'Data Science', 'Data Science', and 'Data Science'. The 'Data Science' workspace is selected, and the 'Overview' tab is active. The console also shows a 'Data Science' workspace with a 'Data Science' workspace and a 'Data Science' workspace.

Administer across multiple types of compute resources from a single interface

Tightly Integrate



Governance



Data
Ingest/Movement



BI Tools



Data Science

Data Fabric

Db2 Andromeda

- _Codename for **next major Db2 LUW release** across all form factors
- _Planned for 2H 2022

- _**Day-and-date** on all three deployment (on-cloud, on-containers, and on-premise)
- _Designed with **cloud-first** in mind

- _Multi-tenancy to allow **multiple logical database**
- _pureScale on cloud
- _**Next evolution of DPF/BLU** for modern warehousing solution

