



IBM Db2 Warehouse

Optimized for always-on analytics
and AI workloads anywhere

Highlights

Empower AI applications with native vector data store and similarity search

Achieve 4 times faster query performance and 34 times lower storage costs¹

Share data and enable a unified view across your analytics and AI estate

Optimize costs through flexible workload scaling

Experience continuous availability and resiliency anywhere

Traditional enterprise data architectures often prove to be complex and expensive and lack the necessary flexibility to fully use data for AI applications. Traditional cloud data warehouses promised a way to drive analytics costs down with pay-as-you-go pricing but led customers to overspending, vendor lock-in and limitations in handling diverse AI and analytics workloads. Companies need a data warehouse capable of handling their mission-critical workloads across both on-premises and cloud environments and integrating seamlessly into their analytics and AI ecosystem.

IBM® Db2® Warehouse is a cloud-native data warehouse that is optimized for always-on, mission-critical workloads. With the native support for vector data type and similarity search, the next generation of Db2 Warehouse delivers advanced capabilities to power AI-driven applications. Through advanced caching techniques and multitiered storage options, it achieves 4 times faster performance and 34 times lower storage costs. It is designed for secure data sharing through native support for open data formats. Db2 Warehouse also integrates with watsonx.data, built on an open data lakehouse architecture, to enable a unified view of your analytics and AI estate. Deploy it as a SaaS on Amazon Web Services or IBM Cloud®, or use the Bring Your Own Cloud (BYOC) model to run it on Microsoft Azure, all as part of your hybrid data management strategy.



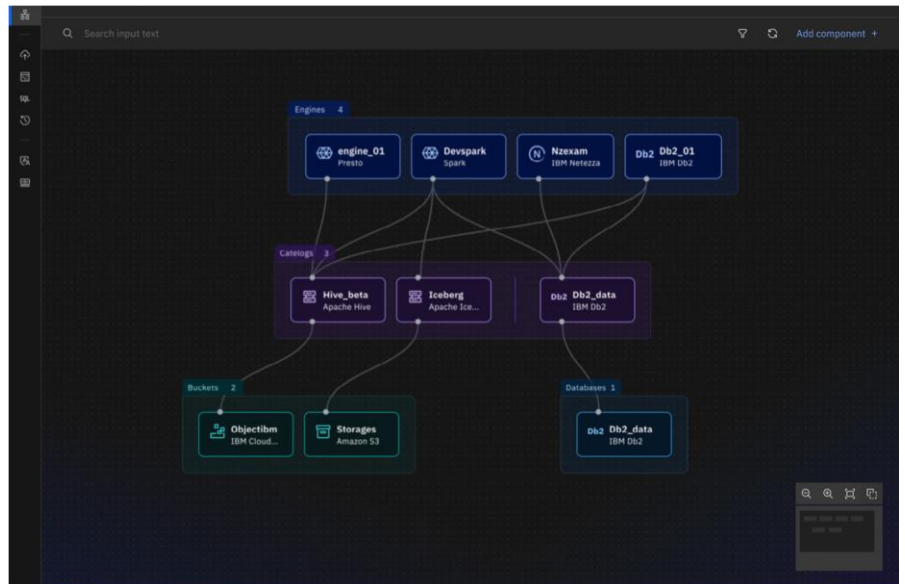


Figure 1. Db2 Warehouse + watsonx.data: existing warehouse workloads across the hybrid cloud with fit-for purpose engines that scale up and down automatically. For example, run resource-intensive ML-model builds in your lakehouse without impacting BI and dashboard workloads in your warehouse.

AI-Ready with Vector Data and Similarity Search

The introduction of a native VECTOR data type enables storage and management of vector embeddings directly within the database. This empowers AI applications such as semantic search, recommendation engines, and Retrieval-Augmented Generation (RAG), allowing developers to use familiar SQL to conduct similarity searches based on vector distances.

With this AI-native functionality, enterprises can build, train, tune, and deploy machine learning models using governed data directly inside Db2, without requiring data movement. Tools like Python, R, and Jupyter Notebooks integrate directly with the database engine for model development.

Achieve 4 times faster query performance and 34 times lower storage costs¹ You might be wondering, “How can I reduce my storage costs but also improve my performance?” The next generation of IBM Db2 Warehouse SaaS brings a host of new capabilities that allow customers to scale their analytics across their entire technology estate while enjoying industry-leading price-performance:

- The solution introduces native support for cloud object storage—based on Amazon Simple Storage Service (Amazon S3)—for Db2 column-organized tables. Coupled with our advanced caching technology, cloud object storage helps clients reduce their storage costs by 34 times and improve performance up to 4 times¹ compared to the current generation service.
- Db2 Warehouse connects dashboards and reports with the power of fast ingest and querying capabilities for real-time insights. The solution speeds up analytics and reporting with a combination of in-memory and column-store data retrieval, maximized CPU processing and data skipping that allows for faster input and output for your data in the cloud.
- Db2 Warehouse is high speed, uses in-memory processing and allows for querying on compressed data. Db2 intelligently caches data that is required for a query execution, leaving the rest on disk.

Share data responsibly and enable a single view across your analytics and AI estate

Db2 Warehouse now gives you the freedom to catalog, ingest and query your data the way you want to with support for open data formats such as Apache Parquet, Apache Avro and the Apache Iceberg open table format. This allows Db2 Warehouse to seamlessly integrate with watsonx.data, an open data lakehouse built to scale AI workloads and provide a common metadata store integrated with multiple query engines. For the first time ever, all your data warehouses and data lakes are stitched together to provide you with a singular view of your entire analytics estate.

- This new architecture also allows you to easily shift workloads and run fit-for-purpose engines such as Presto and Apache Spark in watsonx.data that best meet your workload and price-performance needs.
- Db2 allows for governed data access and the distribution of data among users and third parties without data duplication. [IBM Knowledge Catalog](#) seamlessly integrates with Db2 Warehouse to provide centralized governance, quality and better compliance.
- Seamlessly integrate Db2 Warehouse with IBM business analytics tools such as [IBM Cognos® Analytics](#). And third-party reporting applications such as [Tableau](#), [Microsoft Power BI](#), [Looker](#) and [SAP](#), and ETL tools such as [IBM DataStage®](#) and [Informatica](#) can also be integrated with Db2 Warehouse.
- Use governed data from Db2 Warehouse to build, train, tune and deploy your own ML models. Data scientists and engineers can also use familiar tools such as Python, R and Jupyter Notebooks to analyze and train powerful

machine learning (ML) models on their governed Db2 Warehouse data directly in the database engine, without the need for data movement.

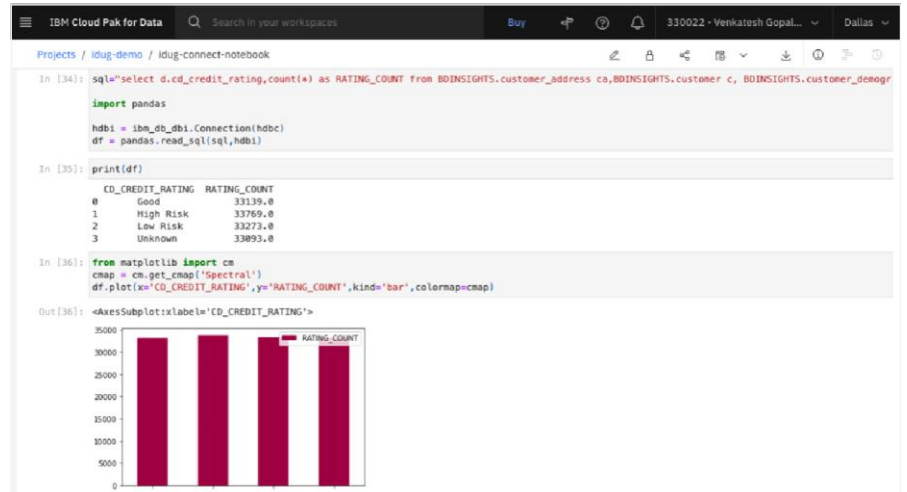


Figure 2. Train powerful ML models with Python directly inside Db2.

Optimize your costs through flexible workload scaling

With Db2 Warehouse, you can take full control of your analytics costs with elastic scale and a cloud-native architecture based on object storage. The solution decouples storage and compute, allowing you to scale compute separately from the storage tier. The vast majority of your data will live in highly scalable cloud object storage, with the option to land your row-organized data in block storage.

Customers can easily scale their data warehouse through the Db2 Warehouse database console. More exact control means paying less without sacrificing performance.

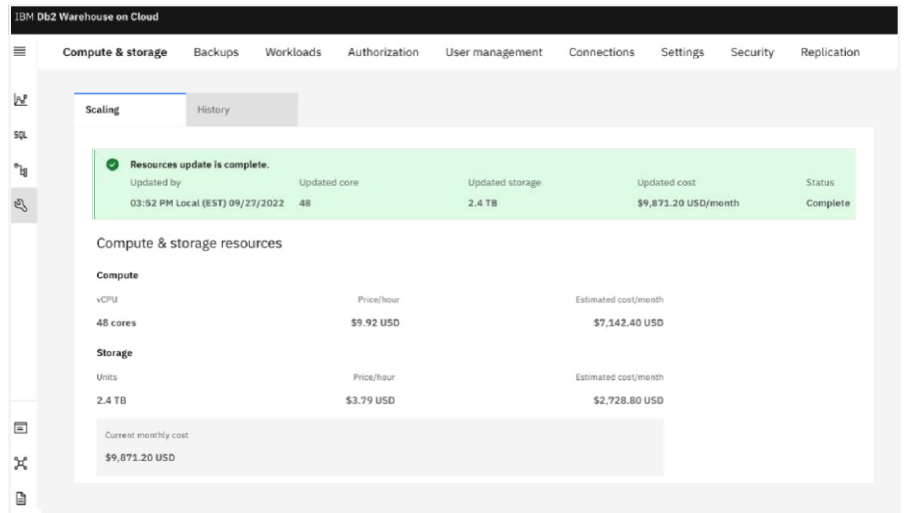


Figure 3. Scale compute and storage based on your resource demands within the warehouse console.

Experience continuous availability and resiliency anywhere

The Db2 Warehouse cloud-native architecture decouples computation and storage and features multiple layers of resiliency with managed computation, highly available storage and cross-cloud replication. If a compute node is in an unhealthy state, the cloud provider's native Kubernetes service—IBM Cloud Kubernetes Service or Amazon Elastic Kubernetes Service—immediately detects the node, removes it from the cluster and delivers a new node from a hot standby pool that is provisioned just in time. Self-service snapshot backup and restore is supplemented by additional disaster recovery Db2 backups that are stored and replicated by the cloud provider's object storage.

Db2 Warehouse includes an adaptive workload management technology that automatically manages resources between concurrent workloads, given user-defined resource targets. Avoid unnecessary downtime or system latency with technology that helps ensure stable and reliable performance when tackling even the most highly concurrent workloads.

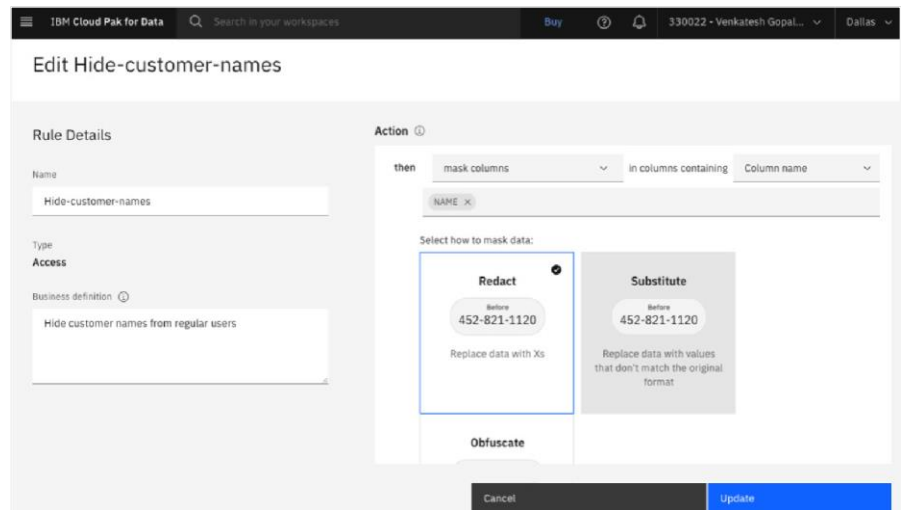


Figure 4. Define database policies to mask and restrict access to data.

Conclusion

The next generation of Db2 Warehouse sets a new standard for enterprise data platforms by combining cloud-native scalability, high performance, AI-readiness, and open architecture. Whether deployed on-premises, in the cloud, or in a hybrid model, it equips enterprises with the tools needed to modernize their data strategy and drive innovation.

Start transforming your analytics estate with IBM Db2 Warehouse today.

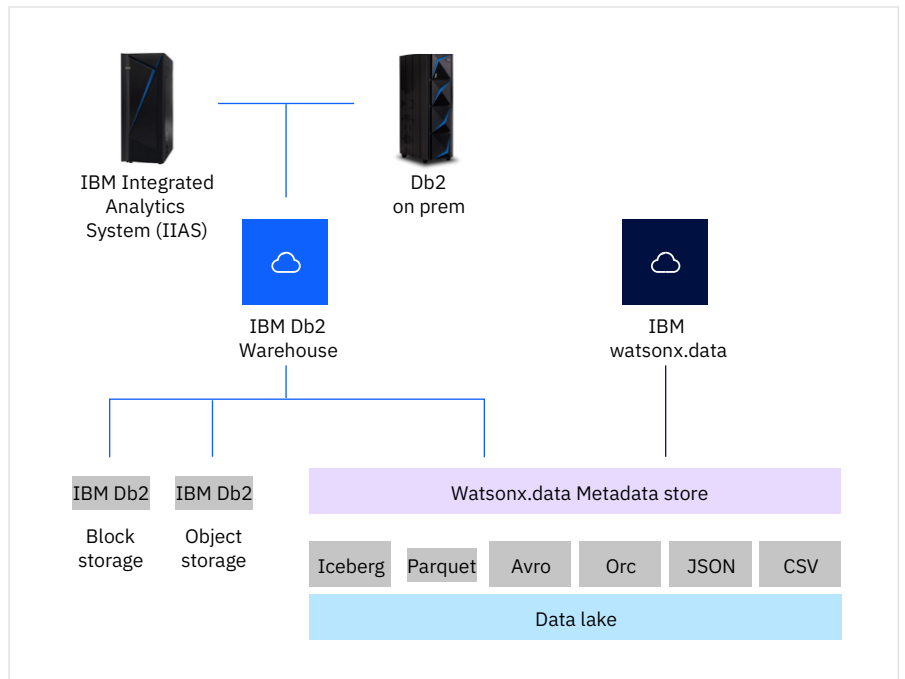


Figure 5. IBM Integrated Analytics System appliance customers and Db2 Warehouse software customers can modernize applications and optimize workloads with Db2 Warehouse, available across hybrid cloud and integrated with watsonx.data.

Why IBM?

IBM has been recognized as a 2022 Gartner Peer Insights Customers' Choice for Cloud Database Management Systems.² Additionally, IBM has been named a leader in the 2023 Forrester Wave: Data Management for Analytics Report.³

For more information

To learn more about IBM Db2 Warehouse, contact your IBM representative or IBM Business Partner, or visit [visit our webpage](#).

Get started with USD 1,000 in free IBM Cloud credit with your instance of Db2 Warehouse Cloud. [Try Db2 Warehouse SaaS today](#).

To learn about Db2 Warehouse pricing, [visit our pricing page](#).

1. Results derived from running the IBM Big Data Insights concurrent query benchmark on two equivalent Db2 Warehouse environments with 24 database partitions on 2 EC2 nodes, each with 48 cores, 768 GB memory and a 25 Gbps network interface. One environment used the new cloud object storage and advanced caching capability, and the other environment did not use the caching capability and was used as a baseline. The test revealed a 4x increase in query speed (213 seconds versus 51 seconds) using the new capability. Reduction of storage costs is derived from price for cloud object storage, which is priced 34x cheaper than SSD-based block storage.
2. IBM is a Leader in 2024 Gartner® Magic Quadrant and also Peer Insights™ Customers' Choice for Cloud Database Management Systems.
3. IBM named a Leader in The Forrester Wave™: Data Management for Analytics, Q1 2023. IBM Blog, 22 March 2023.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.

Statement of Good Security Practices:
No IT system or product should be considered completely secure, and no single product, service or security measure can be completely effective in preventing improper use or access. IBM does not warrant that any systems, products or services are immune from, or will make your enterprise immune from, the malicious or illegal conduct of any party.

© Copyright IBM Corporation 2025. IBM, the IBM logo, IBM Cloud, Cognos, DataStage, and Db2 are trademarks or registered trademarks of International Business Machines Corporation, in the U.S. and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/legal/copyright-trademark.

Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

