

Better together: Next-gen IBM Db2 Warehouse with IBM watsonx.data

Experience a high-performance cloud data warehouse with reduced storage costs and improved performance

The IBM® Db2® Warehouse lets you store and analyze governed data across a variety of sources, formats and analytics environments, while you reduce costs, improve performance and control resource consumption at scale.

The Db2 Warehouse is now seamlessly integrated with the open and governed [IBM® watsonx.data™](#) lakehouse platform, to optimize and scale your analytics and AI workloads across the enterprise.

Achieve 4x faster performance and cut cost by 34x
IBM Db2 Warehouse runs natively on highly resilient, cost-effective cloud object storage. Coupled with advanced caching techniques, your analytics workloads can now run 4x faster while reducing storage cost by 34x.¹

Share data responsibly

Share data in open formats and extract new insights across teams using a single, governed engine that supports integration with your lake.

Scale analytics across the whole enterprise

Db2 Warehouse seamlessly integrates with the watsonx.data service to provide you with a singular view of your entire analytics estate.

Extend warehouse data to your lakehouse

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



Optimize workloads

Optimize costly warehouse workloads with fit-for-purpose engines that scale automatically. Run resource-intensive machine learning (ML) model builds in your watsonx.data lakehouse without impacting business intelligence (BI) and dashboard workloads in your warehouse.

Share data through an open format

Eliminate data silos by sharing Db2 tables with data lakes and lakehouse engines through open table and data formats, such as Apache Iceberg, Parquet, Avro and ORC, and CSV, JSON, and more.

Curate data

Use the watsonx.data lakehouse engine to cleanse and transform data, reducing complexity and time to prepare data for model training. Easily bring lakehouse data back to your warehouse as part of your curated data set, ready for reporting, analytics and AI.

[Learn more](#) →

[Book a meeting](#) →

1. Running IBM Big Data Insights concurrent query benchmark on two equivalent Db2 Warehouse environments; one environment did not use the caching capability and was used as a baseline. A 4x increase in query speed resulted using the new capability. Storage cost derived from price for cloud object storage, which is priced 34x cheaper than SSD-based block storage.

