

Scale AI workloads for all your data, anywhere

IBM watsonx.data is an open, hybrid and governed data store, optimized for all data, analytics and AI workloads



The cost of managing and serving data is rising rapidly as employees at all organizational levels become more comfortable using data. As costs rise, so do scrutiny and demand for a better ROI generated from data, analytics and AI investments. Nontechnical users, such as business stakeholders, increasingly drive the adoption of new analytics tools. And as AI becomes a standard addition to enterprise applications, analytical workload costs will increase even more. It's time for a change.

A new approach to modern data management

A data lakehouse is an emerging architecture that provides a high-performance query engine over low-cost storage and a shared metadata layer. IBM watsonx.data™ is an enterprise-ready data store built on a data lakehouse architecture that enables AI and hybrid cloud analytics workloads such as data engineering, data science and business intelligence, through open-source components with integrated IBM innovation.

50%

Reduction in the cost of your data warehouse through workload optimization across multiple query engines and storage tiers.¹

Highlights

- Use an open and flexible approach to provide a unified view of your data across hybrid cloud environments.
- Deliver price performance through fit-for-purpose query engines such as Presto C++, Presto, Spark and query optimization.
- Improve your AI with quality data through an integrated vector database for retrieval-augmented generation (RAG) and other AI use cases.
- Deploy across any cloud or on-premises environment in minutes with workload portability through Red Hat® OpenShift®.

[Start your free trial](#) →

[Request a demo](#) →

1. When comparing published 2023 list prices normalized for VPC hours of watsonx.data to several major cloud data warehouse vendors. Savings may vary depending on configurations, workloads and vendors.