

IBM Db2 Warehouse on Cloud

A fully-managed, elastic data warehouse for high-performance analytics and AI



Highlights

- Independent, elastic scaling of storage and compute
 - Deployable on multiple cloud providers
 - Columnar-organization and in-memory processing for fast analytics
 - Self-service and geo-replicated disaster recovery backups for data protection
 - Multi-layer resiliency with Kubernetes-managed compute and highly-available block storage
-

IBM® Db2® Warehouse on Cloud is a fully-managed, high-performance, elastic cloud data warehouse. It delivers true elasticity through independent scaling of storage and compute, with a highly-optimized columnar data store, actionable compression and in-memory processing, all working together to supercharge your analytics and AI initiatives. Db2 Warehouse on Cloud can be deployed as a standalone cloud service, or as part of a broader hybrid data management architecture, supporting on-premises data warehouses.





Figure 1: A hybrid architecture allows workloads to move between different types of data warehouses.



Features and benefits of Db2 Warehouse on Cloud

Elastic scaling of storage and compute

Db2 Warehouse on Cloud offers independent scaling of storage and compute, so organizations can customize their data warehouses to meet the needs of their businesses. For example, customers can burst on compute during peak demand, and scale down when demand falls. Users can also expand storage capacity as their data volumes grow.

Customers can scale their data warehouse through the Db2 Warehouse on Cloud web console or API. Scaling compute and storage separately and at a moment's notice is necessary to adapt to unpredictable fluctuations in use. More exact control means paying less without sacrificing performance.

High-performance on complex analytics with BLU Acceleration

Db2 Warehouse on Cloud delivers high performance on complex analytics workloads using [IBM BLU Acceleration](#), a collection of technologies pioneered by IBM Research that features:

- A columnar data store, tuned and optimized for analytics workloads
- High-speed, in-memory processing of columnar data sets
- Querying on compressed data
- The ability to intelligently pull only data required for query execution, leaving the rest on disk

In addition, Db2 Warehouse on Cloud 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 ensures stable and reliable performance when tackling even the most highly concurrent workloads.

High availability and resilience with a Kubernetes-managed architecture

Db2 Warehouse on Cloud's architecture decouples compute and storage and features multiple layers of resiliency by leveraging Kubernetes. If a compute node is in an unhealthy state, the cloud provider's native Kubernetes service (IBM Cloud Kubernetes Service for IBM Cloud, and Amazon EKS for AWS) immediately detects the node, removes it from the cluster, and delivers a new node from a hot standby pool or provisioned just-in-time.

Db2 Warehouse on Cloud's storage layer also utilizes highly redundant, high-performance SSD-backed block storage. Customers can manage self-service snapshot backup and restore through the Db2 Warehouse on Cloud web console. Self-service snapshot backup and restore is supplemented by additional disaster recovery Db2 backups stored and replicated by IBM Cloud Object Storage or AWS S3.

Autonomous and secure cloud service with world-class IBM DevOps support

Db2 Warehouse on Cloud runs on an autonomous platform-as-a-service, and is powered by Db2's autonomous, self-tuning engine. Day-to-day operations, including database monitoring, uptime checks and failovers, are fully-automated. Operations are supplemented by a DevOps team that are on-call to handle unexpected system failure, letting you focus on the data, not the administration.

Built-in AI and geospatial capabilities

Gain deeper business insight and predict future outcomes with AI/ML capabilities. Users can train and run a myriad of machine learning models on their Db2 Warehouse on Cloud data directly in the database engine, without the need for data movement. Algorithms include Association Rules, ANOVA, k-means, Regression, and Naïve Bayes. Db2 Warehouse on Cloud also supports spatial analytics with Esri compatibility, supporting Esri data types such as GML, and supports native Python drivers and native Db2 Python integration into Jupyter Notebooks.

Deployable on multiple cloud providers

Db2 Warehouse on Cloud can currently be deployed on IBM Cloud and Amazon Web Services (AWS) and uses the native building blocks of the cloud provider it's deployed on. For example, when deployed on AWS, Db2 Warehouse on Cloud leverages Amazon EC2 for compute, Amazon Elastic Kubernetes Service (EKS) for container management, Amazon Elastic Block Store (EBS) for storage, Amazon S3 for data lake and data load/unload, and Amazon CloudWatch for monitoring.

Organizations no longer need to experience the stress of cloud vendor lock-in with multiple cloud deployment options. Db2 Warehouse on Cloud brings all of its core capabilities, regardless of the cloud provider it's deployed on. Operating across multiple clouds gives you access to greater data analytic capabilities that may not be found in one, standardized cloud. It also opens the door to be cloud agnostic and access open source technology.

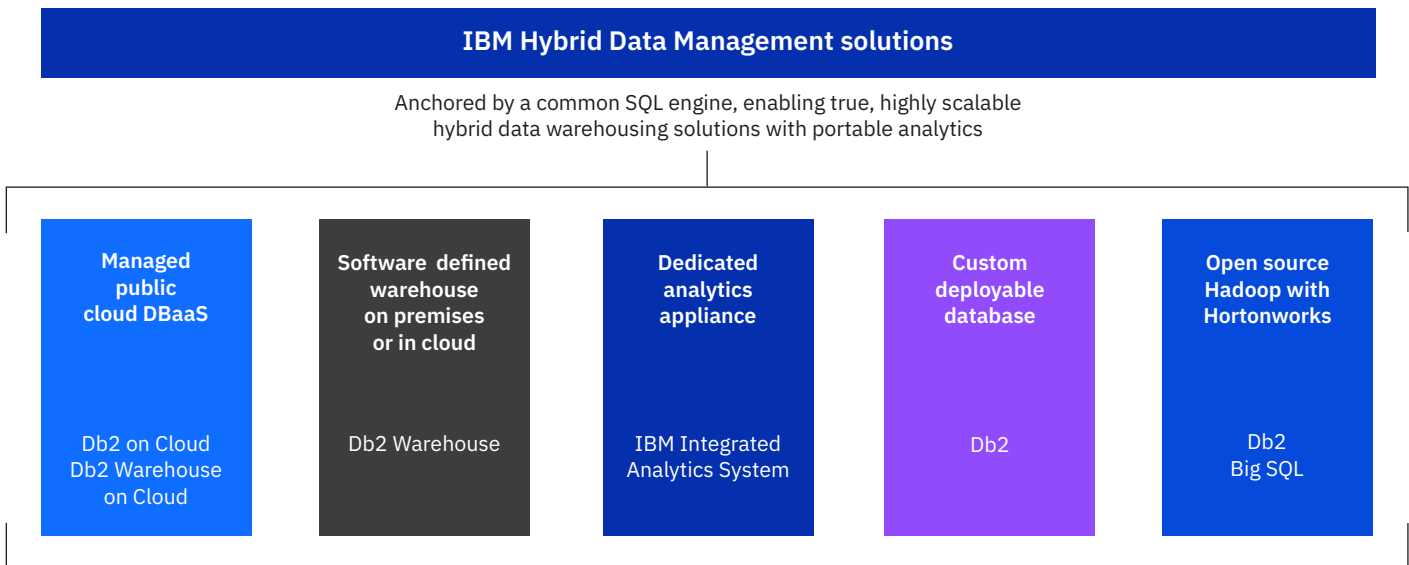


Figure 2: IBM Hybrid Data Management offerings built on the common SQL engine.

A common SQL engine used across the Db2 data management family

With a [common SQL engine](#), clients can architect a true hybrid cloud solution that combines Db2 Warehouse on Cloud with other Db2 form factors, and allows for query portability, operational skill transfer, and data virtualization. The Db2 common SQL engine runs on the same code base, so you can write SQL once and deploy it anywhere. Applications deployed on one Db2 form factor can be easily deployed to any other Db2 form factor, no matter where the database is located, saving you time to focus on your data instead of rewrites.

Use cases

Organizations worldwide use Db2 Warehouse on Cloud to drive their business activities, as part of both hybrid and cloud-native data architectures. Here are a few ways IBM customers are using Db2 Warehouse on Cloud:

Modernize the data warehouse

Reduce time to value with a high-performance, in-memory data warehouse designed for complex analytics workloads and extreme concurrency, while optimizing operational costs with an autonomous, elastic cloud service, designed to grow with the pace of your business.

Transform and analyze mobile, web and the Internet of Things (IoT) data

IoT, mobile, and sensor data are all born in the cloud. Put that data in a cloud data warehouse so you can analyze it in the cloud too, and eliminate the latency incurred when transferring data to an on-premises solution.

Consolidate and integrate data silos and data marts

Using a central cloud data warehouse, instead of many smaller systems provides several benefits due to consolidation. Foremost, data is no longer siloed, meaning that insights can be made with a full breadth of information. Additionally, centralizing the data means data management can be done in a single place, making data management less time-consuming and cheaper.

Bring cloud data to data science

IBM Watson Studio integrates directly with Db2 Warehouse on Cloud, so data scientists and machine learning engineers can use familiar tools like Python, R, and Jupyter Notebooks to analyze and train powerful machine learning models on data that resides in Db2 Warehouse on Cloud.

Service plans

IBM Db2 Warehouse on Cloud offers client three types of flexible and elastic data warehouse configurations to meet their requirements:

		Storage*	Compute
Flex One	IBM Cloud	40 GB—4 TB for user data	6—28 vCPUs
	AWS	40 GB—4 TB for user data	6—28 vCPUs
Flex	IBM Cloud	960 GB—96 TB for user data	16—160 cores; 186 GB memory per 16 cores
	AWS	960 GB—144 TB for user data	14—112 cores; 196GB memory per 14 cores
Flex Performance	IBM Cloud	2.4 TB—96 TB for user data	48—576 cores; 432 GB memory per 24 cores
	AWS	2.4 TB—144 TB for user data	48—576 cores; 432 GB memory per 24 cores

*Storage figures represent disk storage. Data is based on typical compression. Estimated compression is based on historical average of observed data compression rates. Actual client data compression rates and temp space requirements, and resulting data storage availability, are not guaranteed and may vary based on client's specific usage and data characteristics.

For more information and to start a trial at no charge

To learn more about Db2 Warehouse on Cloud and start a trial at no cost, contact your IBM representative or IBM Business Partner, or visit ibm.com/cloud/db2-warehouse-on-cloud.

You can also use \$200 in IBM Cloud credit with your instance of Db2 Warehouse on Cloud¹. [Learn how](#).



© Copyright IBM Corporation 2019

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
September 2019

IBM, the IBM logo, ibm.com, BLU Acceleration, Cognos, DataStage, Db2, and IBM Watson are trademarks of International Business Machines Corp., 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.

Netezza is a registered trademark of IBM International Group B.V., an IBM Company.

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.

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: IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. 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 systems, products and services are designed to be part of a lawful, comprehensive security approach, which will necessarily involve additional perational procedures, and may require other systems, products or services to be most effective. 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.

1 Offer subject to change. Check with your IBM Sales Representative or Business Partner



Please Recycle