IBM Db2 Warehouse

Hybrid data warehousing using a software-defined environment in a private cloud

The evolution of the data warehouse

Managing a large-scale, on-premises data warehouse environments to meet today's growing analytics demands can be complex and expensive. The time is now to consider architecting a hybrid data warehouse that ultimately lowers the cost of analytics, enables unprecedented flexibility and delivers deeper insights.

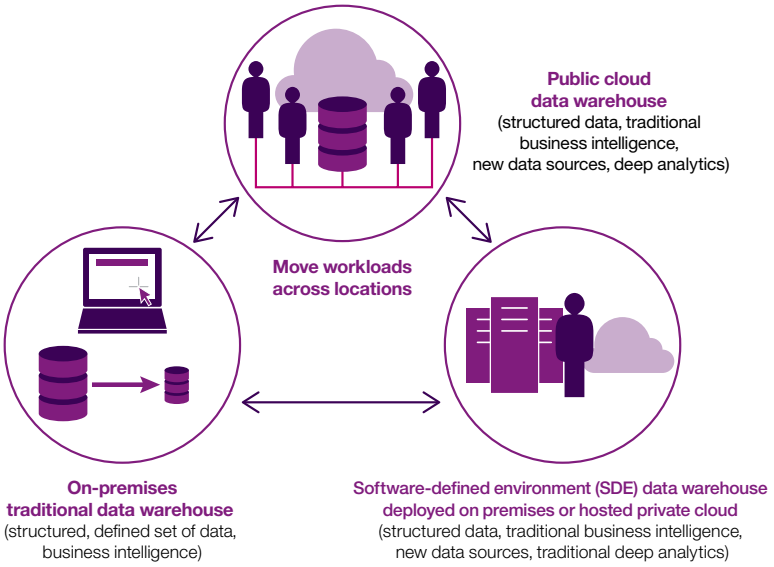

Figure 1: A hybrid data warehouse architecture

A hybrid data warehouse introduces technologies that extend traditional data warehouse capabilities to provide the key functionality required to support new combinations of data, analytics and locations, while addressing the following IT challenges:

- Delivering new analytics services and data sets to meet time-sensitive business initiatives
- Managing escalating costs due to the massive growth of new data sources, analytics capabilities and users
- Achieving data warehouse elasticity and agility for sensitive business data

IBM Db2 Warehouse enables hybrid data warehousing using a software-defined environment

IBM® Db2® Warehouse is a client-managed, preconfigured data warehouse that runs in private clouds, virtual private clouds and other container-supported infrastructures. This data warehouse is designed to provide the ideal solution when you must maintain control of your data, but want cloud-like flexibility. It includes in-memory processing to deliver fast answers to queries, as well as massively parallel processing (MPP) to help you scale out and scale up capabilities as demand grows. For analytics, you can use Db2 Warehouse products to leverage familiar structured query language (SQL), integrated R and Python, or robust in-database analytics—including geospatial analytics.

Db2 Warehouse complements the overall IBM hybrid data warehouse strategy which is designed to provide organizations with the highly flexible architecture that is needed in the dynamic, fast-moving world of big data and cloud computing. Because Db2 Warehouse and Db2 Warehouse on Cloud use a common analytics engine, analytics workloads can be moved across public and private clouds without application changes. Db2 Warehouse technology is compatible with Db2 and IBM PureData® System for Analytics, powered by IBM Netezza® technology, as well as Oracle SQL. This compatibility helps move analytics workloads to Db2 Warehouse or the cloud more easily, depending on the application.

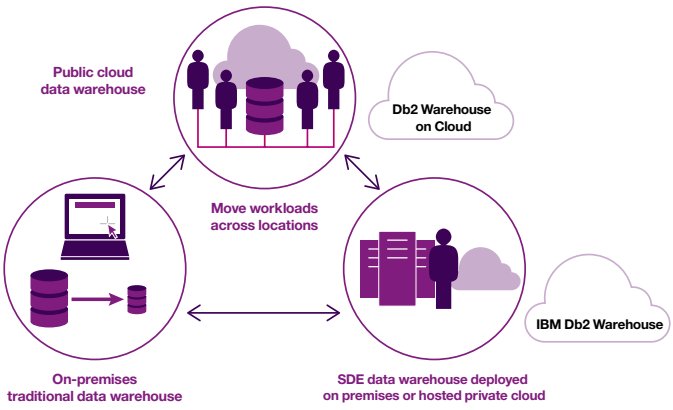

Figure 2: A hybrid data warehouse with Db2 Warehouse products

Db2 Warehouse employs Docker container technology which helps simplify management and reduce deployment times to minutes. It also provides elastic scaling and is designed for ease of updating and upgrading. All of these features are designed to be helpful to IT or cloud administrators. From a user standpoint, Db2 Warehouse helps provide the performance needed to quickly acquire data sets, apply analytics to solve specific business problems, and operationalize insights using the optimal deployment option.

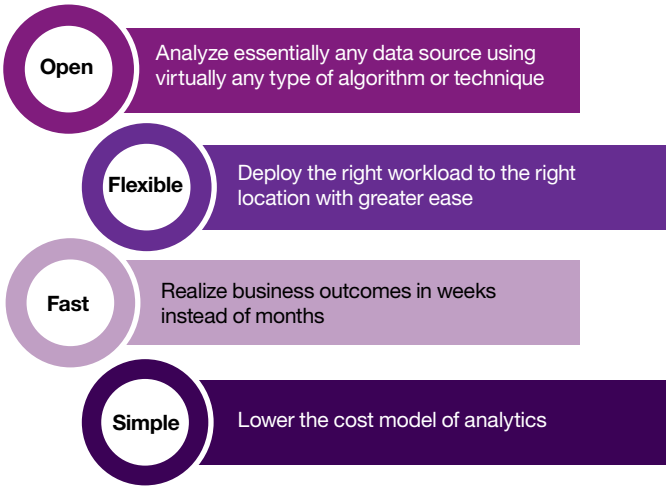

Figure 3: Db2 Warehouse key differentiators

Open

Users demand that IT departments deliver new analytics services, such as R, Python and Spark, as well as support for a variety of data types. Not being able to serve these requests leads to the proliferation of analytics silos and less control of data. Db2 Warehouse helps IT departments gain control by delivering an open solution that is designed to make it easier to process data using virtually any type of algorithm or technique, across a range of data sources. Additionally, you can run Db2 Warehouse on standard hardware supporting Docker installations on Linux, cloud, Apple OSX and Microsoft Windows platforms with minimal prerequisites.

The open nature of Db2 Warehouse allows you to:

- Load a wide range of structured and unstructured data—including geospatial and object storage data—from a variety of sources with greater ease as data is requested.
- Use familiar business intelligence tools, as well as open source R, Python and Spark for in-database processing.
- Connect Esri ArcGIS to perform geospatial analytics.
- Take advantage of the hardware you already employ in your data center.

Flexible

Various teams throughout your organization create analytics solutions using different data sources and tools. To bring all of these solutions together into a single application can require several runtimes, as well as an optimal data flow that may span on-premises and cloud deployments. Additionally, IT departments must efficiently manage workloads to address the latest business needs, such as business-sensitive data and unpredictable demand. Db2 Warehouse technology addresses these challenges with a hybrid data warehouse architecture. Db2 Warehouse provides cloud-like agility and elasticity, while delivering advanced analytics to support the latest programming models and data sources. Because Db2 Warehouse is part of a family of common database technologies, you can write your application once, and move that workload to the right location. Locations include public cloud, private cloud, or on premises—with minimal or no application changes required.

The flexibility of Db2 Warehouse allows you to:

- Choose to run Spark or SQL for analytical processing.
- Move workloads between locations including a public or private cloud, and on-premises data warehouse.
- Leave data where it resides using build-in IBM Fluid Query for federated queries.
- Employ elastic scaling across a wide range of infrastructure resources

Fast

Even organizations with established data warehouses can benefit from a faster method for gaining business outcomes through analytics. With push-button deployment in less than 30 minutes, users can rapidly deliver an optimized, private cloud data warehouse that complements and extends the core on-premises data warehouse. This hybrid architecture jumpstarts new analytics projects with minimal risk. The SDE, private cloud deployment takes advantage of underused resources, with self-service provisioning of the right combination of data and analytics services. By leveraging a single node and MPP architecture for IBM BLU® Acceleration® in-memory processing and Netezza in-database analytics and Spark, users can quickly build and test analytics models against higher volumes of data.

Simple

An SDE is designed to optimize the entire computing infrastructure—including compute, storage and network resources. Additionally, an SDE automatically tailors itself to meet the needs of the required workload. Db2 Warehouse is delivered through Docker container technology and takes advantage of an SDE. For example, it auto-provisions resources to handle changing workload needs. Db2 Warehouse also makes deployment and management more efficient, with elastic scaling, and easy updates and upgrades. Db2 Warehouse provisions a full data warehouse stack, including Spark, in minutes to help you manage the service in your own public or private cloud, while maintaining existing operational and security processes.

Db2 Warehouse is designed for simplicity because of:

- Container technology that eases deployment and management
- System resources that dynamically adjust to satisfy variable workload demands
- Built-in Spark, which means you do not need to install and configure Spark separately
- Embedded high availability and disaster recovery

Is Db2 Warehouse right for you?

Db2 Warehouse can help you address critical requirements, including the need for:

- More data warehouse capacity using a SDE for improved elasticity to continually meet service levels and maximize the use of existing resources, such as commodity hardware
- Resources that can be dynamically provisioned to quickly gain access to the right combination of analytics and data services
- A warehouse or data mart designed to deploy quickly and easily, with little tuning or management
- A cloud strategy built to keep data more directly under business control or on premises due to internal requirements and other mandates
- A cost-effective, high-performance processing engine to gain deeper insights from the massive amounts of data being generated by mobile, web and the Internet of Things (IoT) applications
- A cost-effective alternative to rewriting applications for use with Hadoop, especially when working with structured data and commodity hardware

Get started: Example use cases

The following use cases are meant to inspire you to get started with Db2 Warehouse.

- **Prototyping, development or test ecosystem**. Quickly and more easily test new applications and data sources before production implementation.
- **Departmental or accelerated analytics projects**. Quickly start an analytics service that can meet requirements for handling a range of data sources, advanced analytics and application development.
- **Data warehousing as a service (DWaaS) or hybrid data warehouse**. Partially or fully migrate a subset of applications, data or both to the cloud from an on-premises warehouse.

"For a long time we have been providing traditional, on-premises database services and for approximately five years we have also offered IaaS cloud services. We recently stepped into IBM cloud data and analytic services for the first time and are looking forward to the general availability of [Db2 Warehouse] to utilize in additional customer projects and applications."

— T-Systems

For more information

To learn more about Db2 Warehouse, please contact your IBM representative or IBM Business Partner, or visit this website.

To try Db2 Warehouse, install the Docker engine on the host server and configure POSIX-compliant Clustered File System Storage, such as Global File System 2 (GFS2) and IBM GPFS™.

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: **ibm.com**/financing

IBM Analytics
Route 100
Somers, NY 10589

Produced in the United States of America
July 2017

Please Recycle

IMD14507-USEN-02