



## IBM dashDB Local

*Using a software-defined environment in a private cloud to enable hybrid data warehousing*

### Evolving the data warehouse

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

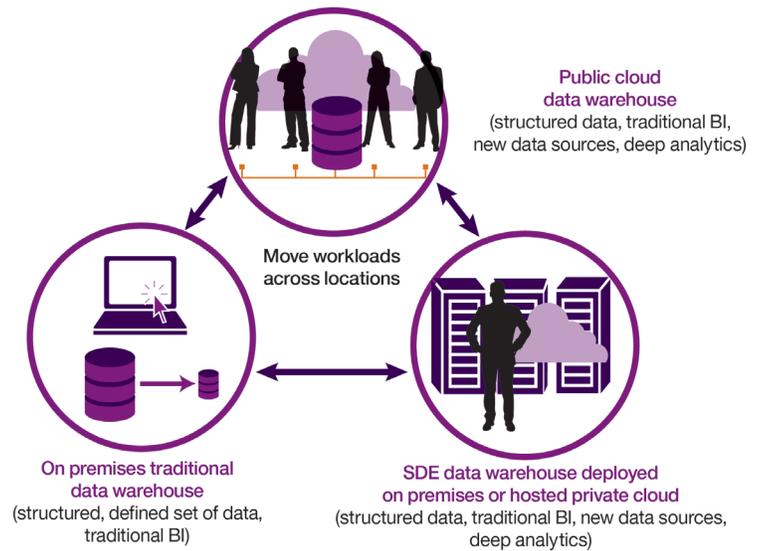


Figure 1: What is hybrid data warehouse architecture?



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

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

### dashDB Local, enabling hybrid data warehousing using a software-defined environment

IBM® dashDB™ Local is a client-managed, preconfigured data warehouse that runs in private clouds, virtual private clouds and other container-supported infrastructures. dashDB Local is designed to provide the ideal solution when you must maintain control of your data, yet 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 dashDB Local to leverage familiar SQL, integrated R and Python, or robust in-database analytics, including geospatial analytics.

As the newest addition to the IBM dashDB product line, dashDB Local complements IBM's overall hybrid data warehouse strategy, providing organizations with the highly flexible architecture that is needed in the dynamic, fast-moving world of big data and the cloud. Due to the common analytics engine between dashDB Local and dashDB managed service, analytic workloads can be moved across public and private clouds without application changes. dashDB technology is highly compatible with IBM DB2® and IBM PureData® System for Analytics (Netezza®), as well as Oracle SQL. This compatibility means that analytic workloads can easily be moved to dashDB or the cloud depending on the application.

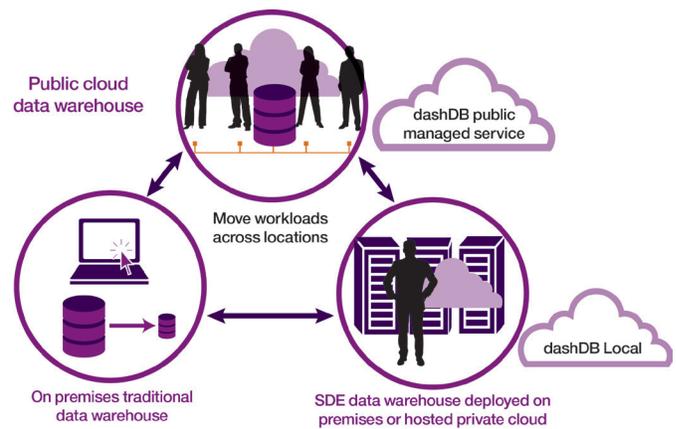


Figure 2: Introducing dashDB Local

dashDB Local employs Docker container technology, which helps simplify management and reduce deployment times to mere minutes. It also provides elastic scaling and easier updates and upgrades. All of these features are designed to be helpful to IT or cloud administrators. From a user standpoint, dashDB Local provides the openness, flexibility and performance needed to quickly cycle through the process of acquiring the right data sets, applying the right analytics to solve specific business problems and operationalizing insights on the right deployment option.

### Open

Users demand that IT departments deliver new analytic services such as R, Python and Spark, as well as support a variety of data types. Not being able to serve these requests leads to proliferation of analytic silos and less control of data. dashDB Local can help IT gain control by delivering an open solution that makes it easier to process data using virtually any type of algorithm or technique, across a wide range of data sources.

dashDB Local is open because you can:

- Load a wide range of structured and unstructured data sources easily as data is requested, including Cloudant® JSON files, geospatial data, object storage and more.
- Analyze larger volumes of data using familiar BI tools, as well as open source R and Python in-database processing, and newer techniques offered through SPARK analytics.
- Connect ESRI ArcGIS to perform geospatial analytics.

### Flexible

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

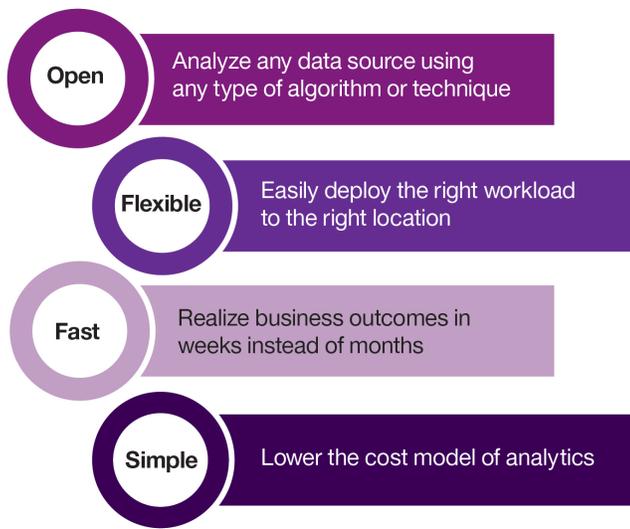


Figure 3: dashDB Local key differentiators

dashDB Local is flexible because you can:

- Easily move workloads between locations.
- Choose to run Spark or SQL for analytical processing.
- Employ elastic scaling for a wide range of infrastructure resources.
- Use federated queries and leave data where it resides.

### Fast

Even an organization with an established data warehouse can benefit from a faster method for gaining business outcomes through analytics. In just under 30 minutes, users can deliver an optimized, private cloud data warehouse that complements and extends the core on-premises data warehouse. This hybrid architecture jumpstarts new analytic projects with minimal risk, because the software-defined environment (SDE) private cloud deployment takes advantage of underused resources and elastic licensing. Diverse data loading and data virtualization options enable users to immediately start exploring volumes of data. By leveraging an MPP architecture for BLU Acceleration® in-memory processing and Netezza in-database analytics, users can quickly build and test analytic models against higher volumes of data.

dashDB Local is fast because of:

- Push-button deployment in less than 30 minutes
- Faster results through in-memory BLU Acceleration, in-database Netezza analytics and Spark
- Out-of-the box transformation of data managed in Cloudant NoSQL database to dashDB Local, including geospatial data
- Transferrable and elastic licensing across deployment locations

### Simple

An SDE is designed to optimize the entire computing infrastructure, including compute, storage and network resources. Additionally, an SDE can automatically tailor itself to meet the needs of the required workload. dashDB Local is delivered via Docker container technology and can take advantage of an SDE. For example, it can auto-provision resources to handle changing workload needs. It also makes deployment and management more efficient, with elastic scaling and ease of updates and upgrades. dashDB Local can provision 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.

dashDB Local is simple 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

---

### Is dashDB Local right for you?

dashDB Local can help you address needs such as these:

- I need more data warehouse capacity. I like the advantage of using an SDE to achieve elasticity to continually meet service levels, and maximize use of existing resources such as commodity hardware.
- I need to dynamically provision resources so users quickly gain access to the right combinations of analytic and data services.
- I need a warehouse or data mart that is designed to be easy and fast to deploy, and requires little tuning or management.
- My company is building a cloud strategy, but we need data to stay more directly under business control or on-premises due to internal requirements and other mandates.
- I need a cost-effective, high-performance processing engine to gain deeper insights from massive amounts of data being generated by mobile, web and IoT applications.
- I don't want to spend millions re-writing applications to work with Hadoop, especially when working with structured data and commodity hardware.

---

### Get started: Example use cases

The following use cases are meant to inspire your thinking about how to get started with dashDB Local.

- **Prototyping, development or test ecosystem:** Quickly and easily test new applications and data sources before production implementation.
- **Departmental or accelerated analytic projects:** Quickly stand up an analytic service that can meet requirements for breadth of data sources, advanced analytics and application development.
- **Data-warehouse-as-a-service or hybrid data warehouse:** If you have an on-premises data warehouse, you can migrate a subset of applications, data or both to the cloud. You have the option to remain hybrid or to fully migrate.

---

*“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 dashDB Local to utilize in additional customer projects and applications.”*

— T-Systems

---

### For more information

To learn more about dashDB Local, please contact your IBM representative or IBM Business Partner, or visit the following website: [ibm.biz/dashDBLocal](http://ibm.biz/dashDBLocal)

To try dashDB Local, [install](#) Docker engine on the host server and configure POSIX-compliant Clustered File System Storage (for example, GFS2, GPFS™).



---

© Copyright IBM Corporation 2016

IBM Analytics  
Route 100  
Somers, NY 10589

Produced in the United States of America  
July 2016

IBM, the IBM logo, ibm.com, dashDB, DB2, PureData, Cloudant, BLU Acceleration, and GPFS 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](http://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. Statements regarding IBM’s future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.



Please Recycle