



---

## Highlights

- Smarter real-time analytics with AI and a fast data store
  - Get more value from streamed data
  - Store and analyze massive data at speed
  - Leverage built-in machine learning capabilities and IBM Watson Studio
  - Avoid vendor lock-in with Apache Parquet open data format
  - Leverage object storage benefits for big data
  - Get started at no charge with the IBM Db2 Event Store developer edition
- 


# IBM Db2 Event Store

*Gain faster, deeper insights for smarter responses to rapidly occurring events with fast data and AI*

Success in business is driven by the ability to uncover trends and patterns quickly so that data-driven decisions can be made at the exact moment when they are most valuable. This is true no matter how fast the data is being created, or how much of it there may be. In fact, this concept has given rise to the term fast data.

Fast data is related to big data, but has some key distinctions that make it perfect for real-time analytics: it takes advantage of very rapid data ingestion and memory optimization for powerful analytics—enabling business insights to be returned in real time or near real time. A variety of use cases benefit from a fast data approach, including those using Internet of Things (IoT) data, digital payment data, bank transactions, retail digital engagement, security alerts, social media data, and weather data.

To get the most out of fast data, companies are also applying AI and machine learning as they develop more advanced event-driven applications—applications that are programmed to respond to events such as a customer clicking on a link in a retail shopping app, information being provided by a machine linked to the Internet of Things, or a communication from a separate program with which the app has been integrated.



Traditionally, many companies have tried to embrace fast data and event-driven apps through streaming analytics that focused on incoming data rather than persistent data. This provided rapid, but limited insight. Previous attempts to include historical data from batch-oriented databases were often too slow when actions (such as ad presentation) needed to occur immediately. For example, an advertisement suggesting a specific shirt might be presented to a customer currently viewing something similar, but important data about that person's price range or color preference, available from multiple searches, may not factor into the ad shown.

IBM Db2® Event Store was designed to alleviate issues such as this, and it includes built-in AI, machine learning and data science capabilities, allowing users to experience the full depth of insight that fast data has to offer. It is a memory-optimized database with extremely high-speed data ingestion and real-time analytics for massive structured data volumes. It transforms fast data analysis by enabling real-time analytics with machine learning on both incoming and stored data—concurrently. It also solves the issue of needing to replicate data across other databases by making ingested data accessible through the open Apache Parquet format. Plus, you can end rewrites and gain enterprise-grade resiliency, scalability and security through use of the common SQL engine.

It supports popular programming languages for building AI-based machine learning applications, including Python, GO, Java, Jupyter notebooks, Watson Studio, Spark ML and other popular choices. And Event Store also runs on Cloud Pak™ for Data, an open, extensible platform that runs on any cloud which provides a single view of all data through data virtualization.

IBM Cloud Paks are newly enabled container platforms natively built on RedHat OpenShift that accelerate your journey to AI. IBM Cloud Pak for Data helps simplify and automate how an organization turns insights into data by providing a unified all in one design that lets users collect, organize, analyze and infuse operational AI throughout their businesses. The capabilities of Db2 Event Store infused inside of IBM Cloud Pak for Data, provide the ability to combine machine learning with real-time enabled databases and streaming frameworks for ingesting analyzing and acting on fast data so that users can realize the value of AI in real time.

Being able to take full advantage of fast data in this way is valuable to fast-paced industries that rely on extreme speed and depth of insight to drive competitive advantage such as finance, retail, manufacturing, utilities, security, and telecom among others.

## Key features of Db2 Event Store

### **Built-in machine learning with IBM Watson Studio**

Built on Apache Spark plus IBM technologies, Db2 Event Store enhances event-driven application development with AI and integrated machine learning. It also includes native integration with IBM Watson Studio, which facilitates building and testing machine learning models. Continual, automated improvement of these models will help develop deeper, more meaningful insights that fuel rapid responses by the application. In addition, IBM Watson Studio helps users learn, create, and collaborate better through tools such as Jupyter Notebooks and RStudio.

### **High-speed data ingestion**

Db2 Event Store is designed with high-speed ingestion technology that's quick enough to capture data at the speeds in which it is generated. It can land 3 million events per second using a 3-node system. That means a single 3-node Event Store cluster is capable of ingesting more than 250 billion events in a day. With this level of speed available, organizations can feel secure in the knowledge that insights that might have otherwise slipped away can now be reliably stored.

### **Real-time analytics**

Db2 Event Store takes advantage of not only streaming data, but also data that has been stored, whether it was gathered just a minute ago or has been stored for a while. This combination promotes better insights by taking advantage of the full amount of data available. The Db2 Event Store query engine is optimized based on IBM's common SQL engine which is meant to protect your data investment. The common SQL engine is part of a comprehensive IBM strategy for flexibility and portability—one that includes application compatibility, strong data integration and flexible licensing. As a result, Db2 Event Store handles analytical workload queries much faster than traditional query engines—allowing insights to be acted upon that much sooner. Plus, you can be confident that once you write a query, it will work across the Db2 family of offerings.

### **Created with open source in mind**

Data is stored using Apache Parquet's columnar data format to enable universal access and easy integration into an open source stack. As a result, vendor lock-in won't be a concern. In addition, since it is built on Apache Spark, organizations can use familiar capabilities like Spark SQL, Spark Streaming, Spark Analytics, and Spark Machine Learning.

### **Highly available**

Db2 Event Store has native high availability built in. This capability allows the cluster to remain operational in the event of a node failure. Queries and ingest can continue uninterrupted even in the event of a machine failure. This is a stark contrast with traditional MPP databases where a single node failure can prevent ingest or query from continuing until the node is restarted. Ultimately, this greater availability means that organizations can rely on Db2 Event Store to be continually available in situations where fast data, fast analytics and deeper knowledge are needed for making smarter decisions.

### **Designed to lower total cost of ownership**

Db2 Event Store is designed to drive down the total cost of ownership through ease of deployment and use. For example, Db2 Event Store is built with Docker containers and managed using Kubernetes so it's easy to install and deploy. Once it has been deployed, the solution can scale compute and storage independently. Db2 Event Store enables you to pay only for the hardware you need at any given time and scale without data movement—something not typically possible with traditional MPP databases. Moreover, the Db2 Event Store integrated environment reduces the need to build and maintain custom environments, which can potentially save costs and improve resource utilization.

## Use cases for Db2 Event Store

### **Risk assessment and crime detection**

Fast data with Db2 Event Store helps improve risk assessment and crime detection by enabling extremely rapid data ingestion from multiple sources, as well as real-time analytics on this data. Using data from disparate sources and multiple locations potentially allows criminals and threats to be identified more quickly. This is true not only for law enforcement, but also for institutions such as banks which are tasked with preventing fraud. The ability to capture and analyze data from multiple sources in real time would allow banks to better identify transaction trends across multiple products and channels. By combining account transactions from across the bank, they could produce a more complete view to better discover risk exposure and fraudulent activity.

### **Smart metering and smarter grids**

The traditional model of collecting one meter reading per month is quickly being replaced by a system where smart meters can collect a reading every 10 or 15 minutes. In other words, the amount of information that used to be collected in 1 year is now being collected every 2–3 hours. On top of that, information on pricing and availability (from energy sources such as coal, solar, wind, and nuclear) is coming in extremely quickly. This presents both a challenge and an opportunity. Using a solution like Db2 Event Store, companies can create fast, high volume integrated platforms for optimizing energy usage, capacity, and billing across a smart grid system. The information that is collected and analyzed better enables organizations to predict peak demands and plan which energy sources to buy from in order to offset spikes in usage. Customers that these businesses serve benefit from the resulting avoidance of brown outs or blackouts.

### **Retail and loyalty programs**

Retail, perhaps more than any other industry, has seen a dramatic increase in data sources that benefit from the fast data movement. This includes in-store beacons that provide hyper-localized geospatial data, and data streaming from mobile applications. Forward-thinking retailers are able to use this data to drive increased loyalty by integrating

streamed payment and couponing events with some of those sources. Deeper insights can be gained by leveraging real time analytics on both the streaming data and the persistent data. By incorporating climate, calendar, mobile, and other forms of streaming data they can measure, refine, and deliver better, more accurate couponing and loyalty experiences that create customers who are thrilled to do business with that retailer again.

### **Manufacturing and industrial monitoring**

The development of the Internet of Things has been a boon for manufacturing and industrial companies. Utilizing IoT data, many companies have already implemented a real-time monitoring framework for automated production lines. Tools such as Db2 Event Store take this one step further by allowing them to ingest a greater volume of readings at higher speed from those machines and compare it more easily against previously stored data. By doing so, they can better determine recent trends in machine performance, not just examine what is happening at the present time. New time-series integrated libraries mean that customers can also call time-series data using special SQL functions. This improvement in reporting means production line faults can be caught before they cause more serious damage, preventative maintenance can be enacted early, and production processes can be optimized leading to greater productivity through better reporting.

### **Healthcare, turbines, connected cars**

Industry or task specific uses of IBM Db2 Event Store are possible based on organizations' individual needs. For example, healthcare providers can take the numerous readings from medical sensors (such as pulse, blood pressure, oxygen levels) which they rely on to detect life-threatening conditions and analyze that data for action in new ways. Energy companies can plan optimal wind turbine usage, and optimize capital expenditures on turbine placement by analyzing factors such as wind data, direction, and speed as well as temperature and humidity. Companies producing connected cars can use data related to weather, traffic, and driving patterns to support automation products. What these examples have in common is that the concept of fast data can be harnessed to help each business perform better.

## Db2 Event Store Editions

	Developer	Enterprise
<b>Ingestion</b>	Ingest and query up to 100,000 inserts per second using a single node	Ingest 1 million inserts per second per node; Scales linearly
<b>Application Development</b>	Develop applications using Python or Scala through notebooks or direct connection to Db2 Event Store. Or use JDBC, ODBC, GO, and all the languages supported by Db2.	Develop applications through Scala, Python or REST using notebooks, or connection to Db2 Event Store. Or use JDBC, ODBC, GO, and all the languages supported by Db2.
<b>Cores</b>	Up to 8 cores	8–56 cores available on Cloud Pak for Data 24–72 cores available with standalone Event Store Additional details at <a href="http://ibm.biz/esmachinesizing">ibm.biz/esmachinesizing</a>
<b>OS support</b>	MacOS Sierra, Ubuntu, Debian, RHEL, CentOS, Fedora, MS Windows 7, 8, 10	RHEL 7.2+
<b>Usage</b>	Development, Experimentation	Production, Testing
<b>Price</b>	No cost for use	No cost for use in pre-production. For production-use pricing, please contact us.
<b>Fast queries through Spark enhancements, indexing and advanced data filtering</b>	Yes	Yes
<b>Includes IBM Watson Studio with machine learning capabilities</b>	Yes	Yes
<b>Writes all data in Apache Parquet format to enable universal access and prevent vendor lock-in</b>	Yes	Yes
<b>Includes embedded console for system monitoring</b>	No	Yes
<b>Write to object storage or use internally provisioned clustered file system</b>	No	Yes
<b>Runs on Cloud Pak for Data</b>	<b>No</b>	<b>Yes</b>

## For more information

Discover the impact that fast data can have on your business today at [ibm.biz/eventstore-pb](http://ibm.biz/eventstore-pb).

### Start using Db2 Event Store at no cost

Start using the developer edition on a single node, or the enterprise edition for pre-production use and testing at no cost by visiting [ibm.biz/try-eventstore-pb](http://ibm.biz/try-eventstore-pb).

### Pricing and contact information

For pricing on full production use of IBM Db2 Event Store please contact your IBM representative or IBM business partner. Contact information can be found at [ibm.biz/contact-eventstore-pb](http://ibm.biz/contact-eventstore-pb).



---

© 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](http://ibm.com), IBM Db2 and BLU Acceleration 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" [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.

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 operational 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.



Please Recycle

---