

# DataStax Enterprise with IBM

A massively scalable, cloud-native NoSQL database built on Apache Cassandra. Deploys based on your enterprise's needs: on-premises, cloud or hybrid.

## Highlights

- A single source for purchasing, deployment, and support of DataStax Enterprise based on Apache Cassandra.
- Highly scalable, highly available with zero points of failure.
- Supports more workload types from graph to search and analytics.
- Designed to meet the availability and performance demands of mobile, web, and IoT apps.
- Leverages IBM's expertise and global footprint.

Apache Cassandra is an open source, distributed NoSQL database that delivers the continuous availability (zero downtime), high performance, and linear scalability that modern applications require, while also offering operational simplicity and effortless replication across data centers and geographies. It is one of the most popular open-source databases.<sup>1</sup> However, Apache Cassandra faces issues like those that plague standard open source software. These issues include scaling, security, lack of advanced data management tools—important features required to run within an enterprise.

DataStax has worked closely with the open source community to build on Apache Cassandra's decade-plus of maturity to solidify its position as the leading database for cloud-native applications. Now IBM, in partnership with DataStax, simplifies the entire enterprise Cassandra experience by providing graph, search, analytics, administration, developer tooling, and monitoring in a single integrated offering.

In addition to this one-stop experience for the procurement, use, and management of Cassandra, users will gain access to IBM's global deployment capabilities and expert support, which includes client consultancy on individual open source database technologies and overall solution architectures.

# The DataStax Enterprise advantage

DataStax Enterprise (DSE), built on Apache Cassandra, is the world's most scalable database, well known for 100% uptime, unmatched low latency, and the ability to handle massive data at planetary scale. DSE delivers enterprise capabilities used by the best internet companies, and trusted by 40% of Fortune 100 and a fifth of Fortune 500 companies. What does the data-driven enterprise get with DSE?

DataStax Enterprise is a flexible, high-performance, and trusted open data stack for the data-driven enterprise, built on NoSQL Apache Cassandra

### Freedom of choice

The flexibility to deploy DataStax Enterprise anywhere, on any platform

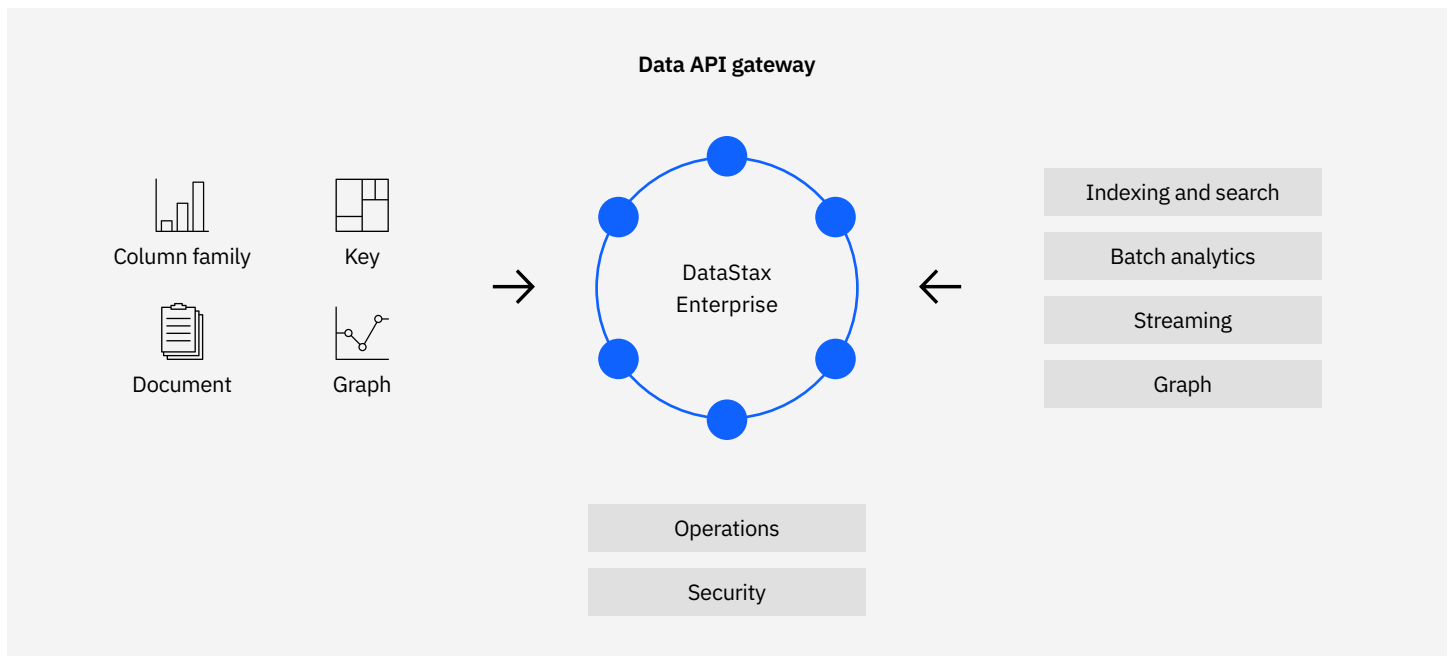
- On-premises
- Bare metal
- Virtual machines
- Any cloud
- Kubernetes

### Superior performance for various workloads<sup>2</sup>

- 50% lower read/write latencies
- 10x better performing graph engine
- 3x faster Spark-based analytical queries
- 4x faster zero copy streaming scaling and node recovery
- 4x faster data transfer than open source cqlsh

### Peace of mind and increased productivity

- Enterprise-grade security
- Easier-to-manage operations
- Reduced learning curve with powerful and modern Stargate APIs including REST, GraphQL, CQL, and schemaless JSON



# Capabilities and benefits

For enterprises that want to self-manage Cassandra clusters, DataStax Enterprise is the only open, multicloud stack built on Apache Cassandra that reduces operational burden, avoids vendor lock-in, and unleashes developer productivity.

## Support for Kubernetes and multiple O/S distributions

- Avoid vendor lock-in with the freedom to deploy on premises, on any cloud, or in a hybrid model.
- Accelerate cloud-native adoption with support for different Kubernetes flavors.
- Greatly simplify cloud-native application development, deployment, and scale-out capabilities with integrated Kubernetes tooling.

## Stargate Data API Gateway

- Reduce the learning curve for getting productive through the abstraction of Cassandra-specific concepts and support for popular API options such as REST, GraphQL, and schemaless JSON.

## Multiple data models

- Gain the flexibility to develop using all styles of NoSQL data models including key-value, wide-column, document, and graph.

## Graph engine

- Explore data relationships fast and at scale with a graph-optimized engine.

## Enhanced search

- Ship search capabilities (including full-text search, faceting, suggester, and geospatial search) with the scalability and resiliency of Cassandra.

## Storage-attached indexing

- Query any column without being restricted to the primary key.
- Stop denormalizing data, and reduce architectural complexity.

## Integrated analytics

- Execute super fast advanced analytics. Inject ML if necessary.
- Support BI and ETL using AlwaysOn SQL.

## Development and migration tools

- Use any popular development language with the appropriate driver.
- Transfer data fast and easily using the bulk data loader.
- Seamlessly migrate data from open source Cassandra.
- Support Apache Kafka streaming with the Kafka connector.

## Enterprise-grade security

- Secure data and protect privacy using encryption, role-based access control, and single sign-on.
- Detect and prevent potential breaches through configurable auditing and log scanning and filtering.

## DataStax OpsCenter

- Easily provision, upgrade, monitor, backup, restore, and manage your DataStax Enterprise clusters using visual management and monitoring tools.

# The added advantage of IBM

IBM brings considerable expertise, strong integration capabilities, and a global presence to improve the value provided by DataStax Enterprise.

## Expertise

IBM has developed considerable experience across thousands of client engagements spanning decades. It shares that experience in two primary ways: consultancy and support. Anyone interested in integrating open source into their data architectures should note that IBM experts are happy to consult on individual open source database technologies or the overall architecture of a business. In this way, an organization can leverage IBM's previous experience to double-check their options and preparedness prior to or even after selecting an open source solution.

Similarly, once a product such as DataStax Enterprise with IBM is purchased, IBM offers their experience in the form of support. 80% of open source support issues are due to lack of product knowledge and environmental compatibilities.<sup>3</sup> This may be because 57% turn to communities for open source support which often fall short in identifying infrastructure and interoperability issues, have slow response times, and lack real-time support.<sup>4</sup> This provides a stark contrast with IBM Multivendor Support Services's (MVS) years of expertise. In a recent Total Economic Impact Study<sup>5</sup> MVS was shown to reduce maintenance and support spending by 25%, lower time spent on hardware support by 20% and reduce time spent on vendor relationship management by 20%.

## Integration

Integration is vital for hybrid data management environments so that businesses can increase efficiencies and drive down costs and wasted time. IBM provides three levels of integration across the various stages of implementation. Foremost, by purchasing open source databases from IBM alongside other hybrid data management technology, an experience that once would have required interaction with multiple vendors can be accomplished easily with one-stop shopping. This also allows the various technologies to be deployed together. IBM also provides one unified architecture for AI so that all data can be brought together in a hybrid environment for self-service analytics with integrated governance. Finally, support can also be delivered through a single source. IBM experts can be leveraged to speak to entire architectures that span multiple offerings whether they happen to be open source or not. Together, these factors provide a seamless environment with a strong vendor relationship.

## Global presence

As a global company, with a presence in over 175 countries and support for 127 languages, IBM understands various geographic challenges and opportunities better than providers who limit themselves to a single market. The advice shared in consultations and support reflects that knowledge. In addition, having data centers around the globe can benefit customers in two key ways. First, locating data nearer to where it will be used reduces latency and helps provide insights more quickly. Second, some regulations require that data is stored in the country of origin. This cannot be done by vendors who are absent from particular markets.

## For more information

To learn more about DataStax Enterprise with IBM, [schedule time with an IBM expert for a no-cost, 30-minute discussion](#). Or visit: [ibm.com/products/datastax-enterprise](https://ibm.com/products/datastax-enterprise)



© Copyright IBM Corporation 2021

IBM Corporation  
New Orchard Road  
Armonk, NY 10504

Produced in the United States of America  
June 2021

IBM and the IBM logo are trademarks of International Business Machines Corporation, 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 [ibm.com/trademark](http://ibm.com/trademark).

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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary. 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.

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.

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.

- 1 <https://db-engines.com/en/ranking>
- 2 Based on IBM internal testing
- 3 2017, Open Source Support Report, Rogue Wave Software
- 4 <https://www.ibm.com/account/reg/us-en/signup?formid=urx-41415>
- 5 <https://www.ibm.com/downloads/cas/GVDZRDZW>