

Modernize Data to Scale AI in Hybrid Cloud

With IBM CloudPak for Data and DataStax Enterprise

Today every organization must become smarter and more agile to intelligently operate and engage the world, resiliently respond to market changes and innovate, and flexibly optimize costs.

To achieve these goals, business applications need to be modernized and address modern use cases. Broadly speaking, business applications fall into two major categories:

- ➔ **The fast lane:** Transactional systems that capture core data that powers the Business and Mission critical functions of the organization.
- ➔ **The deep lane:** Analytical systems that are used for reporting and analysis outside the transactional flows, but on the data gathered by those flows—often from many disparate systems.

Modern “fast lane” applications process data workloads that are used to power e-commerce, IoT, customer loyalty, supply chain management, real-time payments and other modern use cases that require massive scalability and always-on availability. The data resides in various data centers spanning public, private and on-premises clouds, and plays an important role in fueling AI engines that augment human intelligence to make customer facing decisions, and analytical systems that provide business insights. 70 percent of global respondents stated their company is pulling from over 20 different data sources to inform their AI, BI, and analytics systems.¹

Challenges

Having different workloads from different data sources running in data centers that span hybrid clouds can create data silos that prevent business agility and innovation that you want to achieve in the first place.

Further compounding the complexity of these hybrid data landscapes is the fact that the lifespan of that data—the time that it is most relevant and valuable—is drastically shrinking.

Finally, updating old applications in the fast lane can be difficult: your business logic is tightly coupled in a monolith and your data is trapped in legacy systems with no easy migration path to modern, cloud-native alternatives.

Solution

The solution? An agile and resilient cloud-native data platform that can back fast lane microservices, powered by a cloud-native NoSQL data tier. The platform must also host the analytical engines that predict outcomes based on trusted data from various sources.

IBM Cloud Pak® for Data is a fully integrated data and AI platform that enables organizations to accelerate AI-powered transformation by unleashing productivity and reducing complexity. Collect, organize, and analyze data; then infuse AI throughout your business within a collaborative platform experience. Cloud-native by design, IBM Cloud Pak for Data is built on and takes advantage of the underlying resource and infrastructure optimization and management in Red Hat® OpenShift® Container Platform.

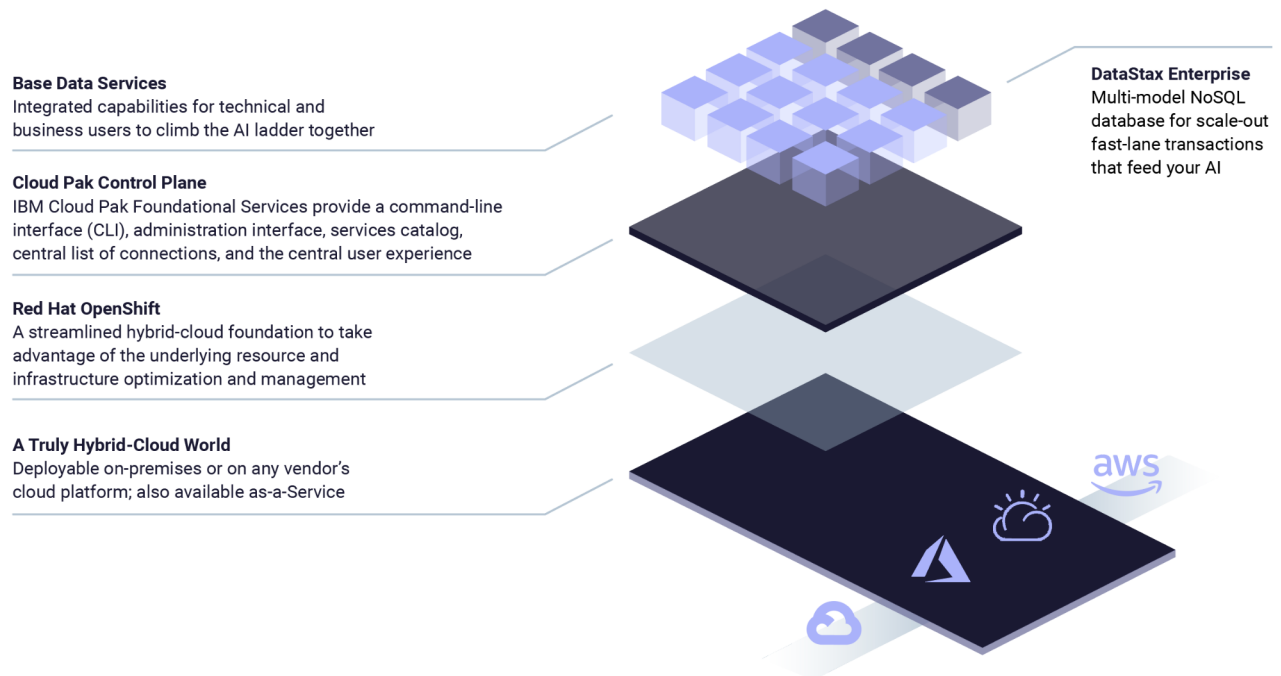
DataStax Enterprise (DSE) is the world’s most powerful, scale-out database for organizations that need to handle modern workloads running in hybrid clouds. Built on the foundation of Apache Cassandra, DSE adds performance, reliability, monitoring, and heightened security.

Certified on Red Hat OpenShift, DSE seamlessly plugs into IBM Cloud Pak for Data as the “fast lane” data tier for NoSQL workloads, and can either run analytical workloads or feed other analytical systems in the Cloud Pak ecosystem.

¹ IBM Global AI Adoption Index 2021 Executive Summary, 2021

Benefits

| | |
|---|--|
| Application and Data Modernization | Run a powerful and distributed data tier alongside your microservices and enable modern use cases at scale in the Cloud Pak ecosystem |
| Superior Modern Workload Performance | Gain blisteringly fast read/write performance for NoSQL workloads and other advanced workloads such as graph, search, and analytical workloads |
| AI Scalability | Automate AI lifecycles for ModelOps by deploying critical capabilities and services on IBM Cloud Pak for Data. |
| Cloud Agnostic | Deploy anywhere on any platform, whether it is in your own data center, in a public cloud, or in a hybrid model |
| Developer Friendly | Prototype and develop on REST/Schemaless JSON/GraphQL, and gain the benefits of the most powerful database with much reduced learning curve |
| Peace of Mind | Gain maximum performance and stability from your NoSQL database backed by the world's foremost Cassandra experts |



Capabilities

IBM CLOUD PAK FOR DATA

Single, unified platform. Bring data management, data governance, data science and AI capabilities together

Designed for Hybrid Cloud. Deploy the platform in almost any environment, whether on-premises, on the cloud or go hybrid.

Built-in governance. Deliver trusted, business-ready data by enforcing policies and rules across your organization.

Pre-built AI and Industry applications. Innovate at speed with solutions for IT operations, customer service, risk and compliance, and financial operations.

Data Science and ML. Drive efficiency in the AI lifecycle to enable data science productivity.

Data Virtualization. Gain centralized access and a single view of data without movement.

Data Fabric. Integrate and engage all your organization's data for better business outcomes.

Data privacy and AI protection. Establish a unified privacy framework to reduce data complexity and manage compliance.

DATASTAX ENTERPRISE

Kubernetes Ready. Run Cassandra workloads on RedHat OpenShift and gain 3x operational productivity.

Modern Data API. Use REST, GraphQL, and Schemaless JSON to access your database and bring applications to market faster.

NoSQL Data Models. Consolidate all NoSQL databases including key-value and document on a single platform

Graph, Search, and Analytics. Grow graph data to billions of graph edges and vertices, and support Solr and Spark workloads.


Storage Attached Indexing. Fast secondary index queries with minimal performance impact.

Painless Migration. Plan and execute migration with tools for unraveling SQL business logic and transferring data seamlessly.

Enterprise Grade Security. Secure data and protect privacy with encryption, security policies, role-based and row-level access controls, and crypto authentication protocols.

Next Steps

- **IBM-DataStax Fireside Chat.** Watch our deep dive into how IBM Cloud Pak for Data and DataStax Enterprise (DSE) enable a Cassandra-based distributed datastore, coupled with best-in-class data science, governance, and privacy.
- **Free Consultation.** Reach out to us to schedule a no-cost consultation to discuss how we can help you modernize your data architecture.
- **Learn More.** Find out more about what IBM and DataStax do together to deliver the data solution for data modernization and scaling AI in the hybrid cloud.



© 2021 DataStax, All Rights Reserved. DataStax, Titan, and TitanDB are registered trademarks of DataStax, Inc. and its subsidiaries in the United States and/or other countries.

Apache, Apache Cassandra, and Cassandra are either registered trademarks or trademarks of the Apache Software Foundation or its subsidiaries in Canada, the United States, and/or other countries.