THE ALWAYS-ON, DISTRIBUTED CLOUD DATABASE BUILT ON APACHE CASSANDRA™ AND DESIGNED FOR HYBRID CLOUD

THE RISE OF THE RIGHT-NOW ECONOMY

DataStax Enterprise (DSE), the always-on, distributed cloud database built on Apache Cassandra and designed for hybrid cloud, powers the Right-Now Enterprise with real-time applications at cloud scale. DSE also gives businesses full data autonomy, allowing them to retain the control and strategic ownership of their data they need to thrive in an increasingly hybrid cloud world.

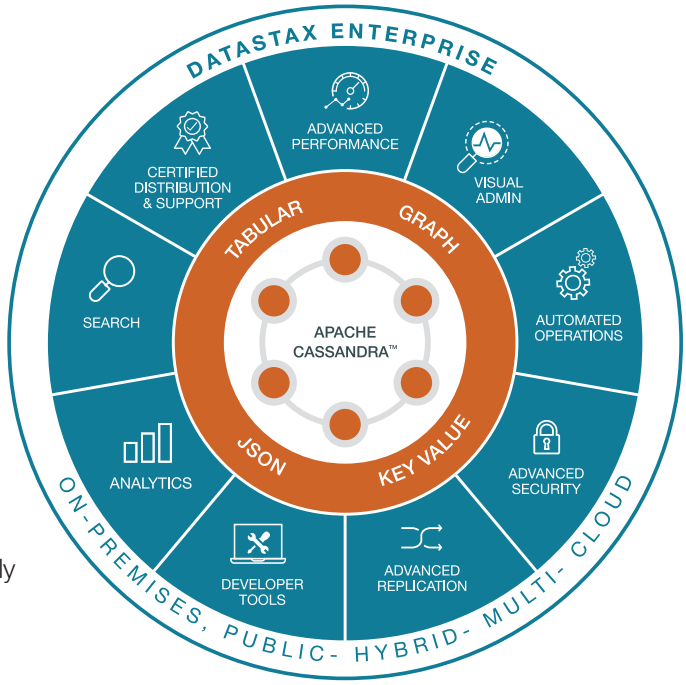

EASY AND FAST DEVELOPMENT

Right-Now Enterprises require easy-to-use tools and flexible data models to build real-time applications quickly and seamlessly leverage data from multiple sources.

- **Multi-Model Development:** Supports multiple data models, including key-value, tabular, JSON, and graph, so you can write data to a single database and access it using a variety of methods based on the needs of your applications.
- **Studio:** Create, explore, and tune CQL, Spark SQL and Graph data with ease using an intuitive, visual interface. Also, context-aware suggestions and validations let you write queries without constantly referring to database schemas.
- **Drivers:** Support for C/C++, C#, Java, Node.js, ODBC, Python, PHP, and Ruby as well as synchronous and asynchronous APIs enables you to easily build distributed, scalable applications.

MIXED WORKLOADS

Every DSE node has integrated indexing, search, analytics, and graph functionality so you can write data once and access it using a variety of workloads or access patterns, all from a single cohesive solution.

- **Advanced Indexing and Search:** Support for sub-string, fuzzy, and full-text search combined with fast, real-time aggregations, faceting, and filtering, allowing you to find data quickly and easily.
- **Robust Transactional Analytics:** Built on a production-certified version of Apache SparkTM, DSE Analytics allows you to easily construct pipelines that interact with both streaming and batch data. DSE Analytics eliminate points of failure, increases performance, and provide full integration with the search and graph capabilities present in DSE. Key capabilities include:
  - DSEFS, which provides a continuously available, easy to deploy, scalable, HDFS-compatible distributed file system for use cases that require a distributed file system for data ingestion, data staging, or state management for Spark Streaming applications.

- AlwaysOn SQL, which ensures analytics have non-stop access to your DSE data via the familiar SQL language as well as BI tools via ODBC and JDBC interfaces.
- Structured Streaming enables simple, efficient, and robust streaming of data from Apache Kafka, file systems, or other sources.
- Spark SQL support on DataStax Studio enables developers to easily create SQL analytic queries and review results via an intelligent, visual query builder
- DSE Analytics Solo, available as an add-on, allows for the creation of a standalone cluster dedicated to operational analytics, ensuring predictable application performance.

INDUSTRY-LEADING PERFORMANCE

Right-Now Economy applications demand guaranteed and immediate responsiveness, no matter the surge in users or workloads. DSE provides:

- **Advanced Performance:** 2x more throughput for both reads and writes and half the latency over Apache Cassandra.
- **Bulk Loader:** Move data into and out of DSE up to 4x faster than other data loading utilities.
- **Faster Analytical Queries:** Retrieve data from DSE 3x faster than open-source Spark reading from open-source Cassandra.
- **Advanced Replication:** Builds upon the multi-data center support in Cassandra to facilitate “hub and spoke” replication requirements for asynchronous connectivity issues common with remote deployments.

OPERATIONAL SIMPLICITY

DSE enables even novice DBAs and DevOps personnel to manage operations like seasoned professionals, reducing operational overhead and increasing productivity.

- **Visual Management and Monitoring:** Easily and proactively provision, monitor, backup, and manage your DSE clusters via DataStax OpsCenter.
- **NodeSync:** Eliminates manual repair operations by keeping data synchronized in DSE clusters in a transparent and continuous fashion.
- **Upgrade Service:** Provides error-free patch upgrades—automatically notifies that an upgrade is available, downloads the software you need, and applies it to a cluster with zero downtime.
- **TrafficControl:** Delivers advanced resiliency that ensures DSE nodes stay online under extreme workloads.
- **Tiered Storage:** Intelligently moves older and less frequently accessed data to cheaper tiers of conventional storage, allowing you to achieve performance and cost optimizations.

END-TO-END SECURITY

DSE ensures business-critical data is protected and helps meet compliance requirements like HIPAA, PCI, and SOX, and the European Union's GDPR by providing:

- **Unified Authentication and Access Control:** Easily provision and manage user permissions via role assignments and a variety of authentication sources, including LDAP, Active Directory, and Kerberos in the same cluster.
- **End-to-End Encryption:** Protect your data on disk, in transit between nodes, and between client and server using advanced integrated encryption capabilities and support for the KMIP encryption key management protocol.
- **Data Auditing:** Track and log all user activity on DSE clusters so you can detect any suspicious activity early on and meet industry compliance requirements.

To learn more about the different subscription models Datastax has to offer, please click here.

Last Rev: MAY2019

ABOUT DATASTAX

DataStax delivers the always-on, active everywhere distributed hybrid cloud database built on Apache Cassandra™. The foundation for personalized, real-time applications at scale, DataStax Enterprise makes it easy for enterprises to exploit hybrid and multi-cloud environments via a seamless data layer that eliminates the issues that typically come with deploying applications across multiple on-premises data centers and/or multiple public clouds.

Our product also gives businesses full data visibility, portability, and control, allowing them to retain strategic ownership of their most valuable asset in a hybrid/multi cloud world. We help many of the world's leading brands across industries transform their businesses through an enterprise data layer that eliminates data silos and cloud vendor lock-in while powering modern, mission-critical applications. For more information, visit www.DataStax.com and follow us on Twitter @DataStax.

DATASTAX