

Lower Cost of Ownership and Get Faster Insights from Big Data

Enterprise-Class Converged Database Platform for Big Data

With a unique converged Big Data vision, Esgyn™ empowers global IT organizations to deliver value to their business from a single Big Data platform for all kinds of workloads (HTAP/Translytics). Enterprises can now enter the emerging era of Operational Big Data to migrate operational workloads to EsgynDB™ / Apache® Hadoop® and enjoy open source economics without extensive application rewrite.

- Transactions to Analytics.** Run transactions, reporting, BI and analytics on the same Big Data Platform
- Built on Apache Trafodion.** EsgynDB is built on Apache Trafodion (incubating) and Hadoop, which means you can enjoy lower total cost of ownership
- SQL Power with NoSQL flexibility.** Get the best of both worlds – flexibility and scalability of NoSQL with ease of use and streamlined application access of RDBMS
- Minimal Data Movement.** Run transactions, operational reporting, BI and analytical workloads on structured, unstructured and semi-structured data so that you can get the advantage of minimal data movement and increased performance.

The Ultimate Query Only EsgynDB can support querying and joining data from multiple storage engines and multi-structured data.

Select * from struct_table, unstruct_table, semi_struct_table, external_source

The Most Capable EsgynDB is the only SQL-on-Hadoop solution to run the entire TPC-DS benchmark.

Technical Features

ANSI SQL
ODBC, JDBC, .NET, Hibernate
Distributed ACID Transactions
HBase, Hive, ORC, Lucene, Spark
Enterprise Class, High Availability
Extensible Data Pipeline with UDFs
Security (Sentry), Manageability

Key Benefits

Reduce server footprint
Scale at a lower cost – petabytes
Use existing SQL resources and tools
Empower Data Scientists and Analysts
Leverage existing BI tools such as Tableau
Flexibility of On-premises or Cloud deployment



“Use of Esgyn technology has provided Webroot with an ability to improve our threat intelligence. EsgynDB accelerated our time to bring new solutions to market due to simplicity of a standard ANSI SQL interface. With EsgynDB we now have the cloud-based architecture we wanted to deploy new and additional workloads quickly and easily.”

Kristie Shetye, Sr. Manager, Software Development, Webroot

Use Cases

Operational Data Store (ODS)
 Self-Serve BI on Hadoop/Data Lake
 Industrial Internet of Things (IoT/IIoT)
 Offload ETL/RDBMS/Reporting
 Offload EDW to Hadoop
 Text Mining/Machine Learning

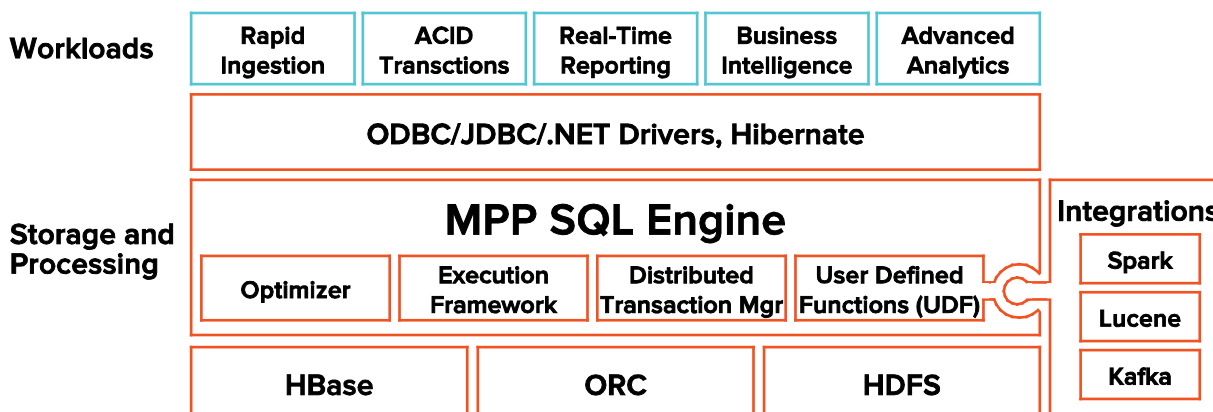
Choice of Deployment

On-Premises (Software)
 Pre-configured Appliance (4-node, 8-node)
 Cloud (AWS, Microsoft Azure or private cloud)
 Big Data Appliance (SuperMicro)

Why EsgynDB?

- Most mature and capable MPP SQL Engine
- Sophisticated Optimizer
- Parallel Execution Framework
- Linear Scalability
- Eliminates data silos
- 170+ patents
- Certified on Cloudera CDH 5.4, 5.5, 5.6, 5.7 and HortonWorks HDP 2.3, 2.4
- Supported on CentOS 6.5-6.7, 7.2 and SUSE SLES 11 Patch 3
- Origins of HP and Tandem

Architecture



About Esgyn

Esgyn is the leader in Big Data solutions that empower global enterprises to realize the potential of Big Data. With the industry’s most mature and scalable SQL Database for Big Data, Esgyn helps enterprises cope with ever increasing data management needs.