ot opentext™

OpenText Analytics Database (Vertica)

A deploy-anywhere SQL database designed for
advanced analytics, speed, and elasticity

Highlights

- **Standards-based analytics:** Speaks SQL, R, and Python so everyone in your organization has access to analytics

- **Advanced analytics:** Includes 700+ analytical functions, including time-series, geospatial, and embedded ML

- **Performance and speed:** Massively parallel and proven to handle the world's biggest workloads and highest number of queries

- **Efficient:** Through proprietary design, it uses fewer nodes and less storage than any other database

- **Flexible deployment:** SaaS, public clouds, private clouds, Kubernetes, and bare metal deployments offer analytics when and where you need it

- **Ecosystem-friendly:** Connects to ETL and BI tools and hundreds of middleware applications

- **Data warehouse/lakehouse:** Store data in our optimized repository or leverage your own Parquet files for analytics

- **Secure:** Encrypts data in the database and data in-motion, as well as a host of other features

- **Eon mode:** Enables cloud-native elasticity with compute-storage separation for fast, scalable analytics and cost efficiency

- **Data visualization:** Allows users to easily create interactive dashboards, helping uncover insights from large datasets quickly.

With the continued growth of data volumes and citizen data scientists' broader use of analytics, many companies are re-examining their systems to meet modern demands. Analytics is rapidly evolving. New data from IoT sources, social media, Weblogs and data streams, gas and electrical grids, and mobile networks is being collected in massive data sets. This gives organizations a new opportunity to become truly data-driven, provided they can manage the new data growth and discover the patterns and trends that can lead to both business opportunities and repeat business from their customers.

Product overview

OpenText™ Analytics Database can help you and your team modernize your data warehouse, deploy analytics in an on-premises/hybrid cloud environment, and democratize data and analytics to provide further access. In addition, OpenText can unify how your company powers its analytics by offering an open, scalable and elastic database with many powerful features.

Amazon Web Services · Microsoft Azure · Google Cloud

Public cloud

On-premises
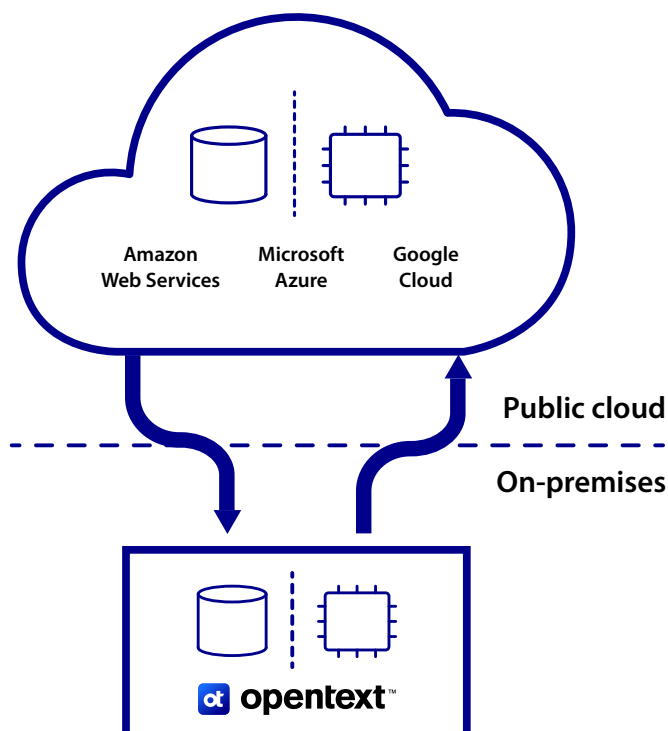
ot opentext™

OpenText can help unify the flow between on-premises and the public cloud.

Product highlights

Eon Mode

OpenText Analytics Database in Eon Mode is a cloud-native architecture that separates compute and storage, providing scalable, cost-effective analytics. This flexible model enables organizations to independently scale compute resources to meet changing workload demands, optimizing both performance and cost. Eon Mode leverages low-cost storage for data while dynamically allocating compute power as needed, ensuring efficient resource utilization. With high availability, workload isolation, and support for multi-cloud and hybrid environments, the Eon Mode ensures fast, reliable, and elastic analytics to support the needs of modern data-driven enterprises.

Standards-based analytics

OpenText Analytics Database offers complete SQL-based analytical functions that make the best use of all the memory and all the resources of the OpenText cluster. SQL is our first and most used language, making it friendly to analysts and data scientists alike. While OpenText offers some important differences behind the scenes, such as column data storage and massively parallel processing, users won't have to write special queries or consider data structure at all. Instead, you run your SQL, and it runs fast.

Advanced analytics

With OpenText Analytics Database, you can finally get machine learning into production. OpenText supports cluster optimized ML (machine learning) algorithms, R, and Python. Data scientists and analysts can build their models using their preferred tools and languages, then leverage OpenText to power them on bigger data sets. In-database machine learning addresses every step in the ML process.

More experienced teams can import models built and trained in other platforms and languages—like TensorFlow, Spark, Python, and SPSS—via the PMML (predictive model markup language) format. With PMML model export, you can export models created in OpenText for scoring in other systems, such as edge nodes for IoT use cases. Data science and data engineering teams can finally operationalize machine learning while using their tools to fully leverage their data at scale.

Performance and speed

Exceptional analytical performance requires efficient data loading, lightning-fast query speed, and high concurrency. OpenText Analytics Database is built from the ground up to blast you through all three. It offers superior load throughput, blazingly fast analytical speed, and high concurrency for multi-user querying. You get speedy query results, reducing query times from hours to minutes or minutes to seconds—something outdated row-store technologies cannot achieve. Its modern architecture gives you the flexibility and scalability to ramp up as workloads increase. OpenText Analytics Database can even isolate workloads to serve multiple departments without duplicating data. OpenText's sub-clustering capability allows workload isolation, eliminating resource competition between BI and data science projects. End-to-end machine learning management—from data prep to deployment. No need to download and install separate packages.

Efficiency

OpenText Analytics Database offers better compression and impressive optimizations. As a result, you will require less hardware and storage than comparable data analytics solutions. In addition, our optimized architecture means that queries run 10 to 50 times faster than on most platforms, while you can store between 10 to 30 times more data per server.
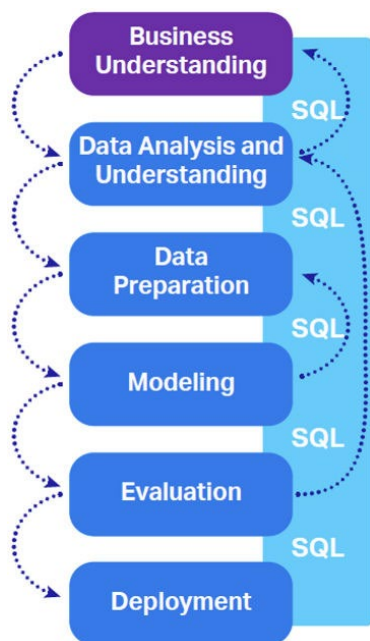
OpenText Analytics Database built-in machine learning process flow.

Efficiency comes in the form of tools, too. OpenText Analytics Database Designer improves operations and reduces resource consumption while improving your team's projection designs for consistently faster queries. It analyzes your logical schema, sample data, and your sample queries. OpenText Analytics Database Designer creates a physical schema design (a set of projections) that you can deploy automatically or manually. It can be used by anyone, even users without specialized database knowledge.

Hybrid cloud

OpenText Analytics Database helps you avoid getting locked into a single cloud vendor—you can use the tools of your choice and take full advantage of the underlying infrastructure you already have in place. It offers portability across multi-cloud, on-premises environments, and data lakes. OpenText runs on popular public clouds, including Google Cloud Platform (GCP), Azure, AWS, Alibaba, and VMware clouds. However, it can run in a similar configuration on premises while leveraging a range of local object stores, such as Apache Hadoop HDFS for communal storage, MinIO, Vast, Dell ECS, Scality, H3C, and Pure Storage. It creates a tough, flexible platform for running a company's analytical and computing workloads, because applications can run simultaneously on multiple environments in a hybrid cloud infrastructure. OpenText can use public clouds and private data centers, and offers the flexibility to switch at an instant.

Ecosystem-friendly

Every release of OpenText Analytics Database is certified and tested with visualization and ETL tools. It supports popular SQL and Java Database Connectivity (JDBC)/Open Database Connectivity (ODBC). You can preserve years of investment and training in these technologies because all popular SQL programming tools and languages work seamlessly. Leading BI and visualization tools are tightly integrated, such as Tableau and MicroStrategy, as well as popular ETL tools like Informatica, Talend, Pentaho, and more.

Data warehouse/Data lakehouse

OpenText Analytics Database can analyze data wherever it resides—HDFS or Cloud Object Storage—and in all popular formats—ORC, Parquet, JSON, or ROS (native Vertica). You do not need to port data simply for analytics from one location to another. Instead, you save time and realize data insights sooner with the data in place. OpenText also analyzes complex data types like Maps and Arrays, Structs in Parquet on S3, HDFS, open SQL-based analytics, and new use cases.
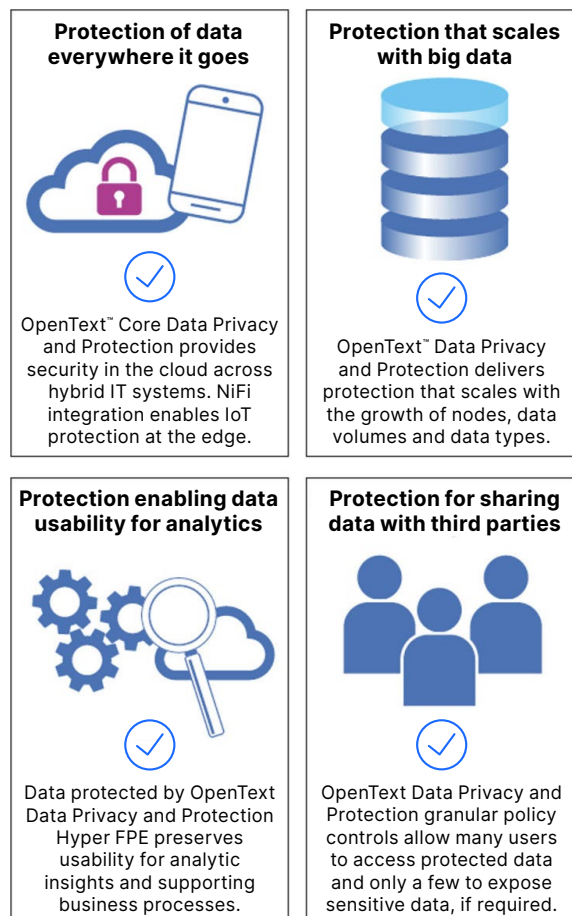
Full and complete analytics

OpenText Analytics Database offers more built-in analytics features than any other data analytics platform:

- Many powerful SQL commands not found on lesser databases, including time-series and geospatial
- Simple SQL execution of predictive analytics (ML/AI) gives analysts the key to machine learning
- OpenText supports Python and let's you use the OpenText cluster to speed through your Python code
- Prepare data with functions for normalization, outlier detection, sampling, imbalanced data processing, missing value imputation, and more
- Create, train, and test advanced machine learning models on massive data sets
- Evaluate model-level statistics including ROC tables and confusion matrices
- Support for familiar programming languages lets you develop userdefined extensions (UDx) with C++, Java, zPython, or R.
- Built-in machine learning algorithms support classification, clustering, and predictive applications. These include linear regression, logistics regression, k-means, naïve Bayes, support vector machines and random forest.

Security

OpenText Analytics Database provides end-to-end security with support for industry-standard protocols and partner solutions such as LDAP, Kerberos, TLS, FIPS 140-2, AWS IAM, and Apache Sentry. OpenText uses a layered security model and provides multiple security authentication authorization mechanisms. Authentication and access is controlled by passwords stored with SHA, LDAP, Kerberos, and SSL certificates. It allows fine grained access control to row and column data, including column masking. It keeps an audit trail, natively exported to other security domains for analysis and persistence.

OpenText Analytics Database's integration with OpenText™ Data Privacy and Protection format preserving encryption technology ensures end-to-end data protection, securing data in use, in motion, and at rest. OpenText Data Privacy and Protection adds a data-centric layer to OpenText's layered security model, facilitating compliance with new and emerging data privacy regulations such as CCPA and GDPR. OpenText Data Privacy and Protection encryption reduces the risk of breach by securing sensitive data while preserving the data's length and character set. Encryption protects structured data such as tax ID, name, address, GPS location, IP address, date of birth, and salary—all typical personal and protected health information types. The decryption of protected data is controlled by policy, including requiring LDAP authentication and authorization if needed.

Protection of data everywhere it goes

OpenText™ Core Data Privacy and Protection provides security in the cloud across hybrid IT systems. NiFi integration enables IoT protection at the edge.

Protection that scales with big data

OpenText™ Data Privacy and Protection delivers protection that scales with the growth of nodes, data volumes and data types.

Protection enabling data usability for analytics

Data protected by OpenText Data Privacy and Protection Hyper FPE preserves usability for analytic insights and supporting business processes.

Protection for sharing data with third parties

OpenText Data Privacy and Protection granular policy controls allow many users to access protected data and only a few to expose sensitive data, if required.

OpenText Analytics Database's integration with OpenText Data Privacy and Protection for enhanced security.

Data visualization

Data visualization in OpenText Analytics Database enables users to easily explore and analyze complex data through intuitive, interactive charts and dashboards. OpenText Analytics Database integrates seamlessly with popular visualization tools like Tableau, Power BI, and Qlik, as well as OpenText Intelligence (Magellan), allowing users to create powerful visualizations that

bring data insights to life. By leveraging high-performance advanced analytics, users can visualize large datasets with near-instant query response times, making it easier to uncover trends, patterns, and anomalies. This empowers users to make data-driven recommendations quickly and influence informed decision-making, enhancing both operational efficiency and strategic planning.

Choose OpenText-as-a-Service, or customer-managed OpenText Analytics Database

OpenText™ Core Analytics Database is our data analytics SaaS offering, which provides all the functionality of OpenText Analytics Database's OpenText Analytics Database (Vertica), available as software as a service on Amazon Web Services. It runs in your own Amazon cloud account, and OpenText experts will help you with onboarding. All of your data analysts and data scientists can tap the power of OpenText advanced analytics and machine learning without worrying about scaling, upgrades, monitoring, or support.

The OpenText Analytics Database platform is our customer-managed offering—the revolutionary column-store, relational database designed from the ground up for analyzing massive data sets, whether on premises, on Hadoop, in the clouds, in any combination of those.

OpenText's multi-parallel processing (MPP) and shared-nothing architecture remains unrivaled for speed and concurrency, and offers in-database machine learning capability for teams seeking advanced AI solutions without compromising workloads for critical business intelligence SLAs.

Try OpenText Core Analytics Database or the OpenText Analytics Database Community Edition, for free

Your data analytics needs are unique, and always changing. Your analytics data warehouse should be flexible enough to meet your demands today, while helping you see into your business's future.

Try OpenText Core Analytics Database, and make your concept a reality.
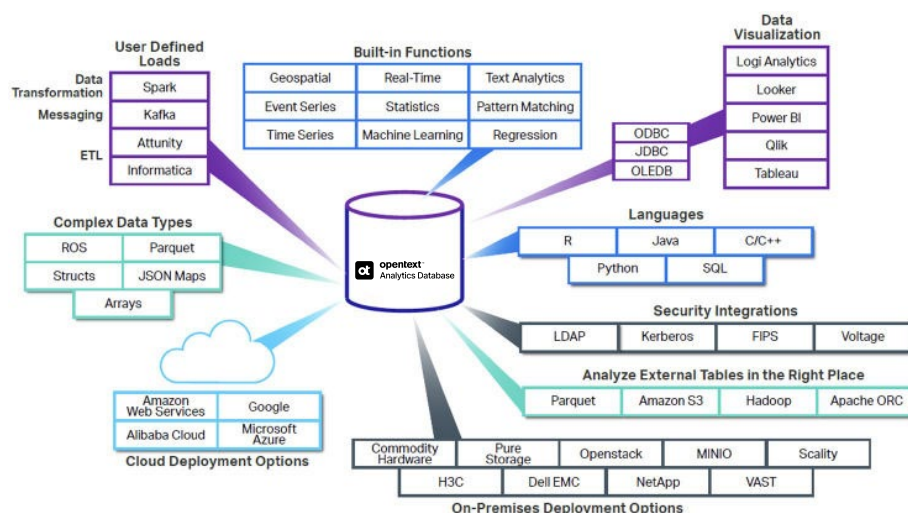
OpenText Analytics Database's open architecture and rich ecosystem.

opentext™