

Vertica Advanced Analytics Platform 9.1

Vertica 9.1 features Eon Mode cloud-optimized separation of compute and storage architecture, advancements in in-database Machine Learning, and security and performance enhancements—now available via a convenient AWS usage-based subscription

Product Highlights

Vertica performs queries 10 to 50* times faster than traditional analytical databases, at a fraction of the cost and using a fraction of the compute and storage resources. In its latest release 9.1, Vertica in Eon Mode enables organizations to optimize infrastructure costs and simplify operations for their Vertica cloud deployments on Amazon Web Services (AWS) by separating compute resources from data storage. Vertica users in the AWS ecosystem can now load high volumes of data into AWS S3 for cost-effective storage and use Vertica's query-optimized ANSI SQL analytics engine for advanced analytics to capitalize on cloud.

Vertica 9.1 provides end-to-end protection of Vertica data at rest, in motion and in use via tighter integration with Voltage SecureData. Vertica 9.1 also includes new in-database Machine Learning capabilities—including Principal Component Analysis (PCA), a popular choice for feature reduction and model efficiency and accuracy. With this release, organizations now have the broadest choice of deployment options in the industry for on-premise, natively on Hadoop, hybrid, and cloud workloads—all delivered via the unified Vertica analytics database.

Key Features

Vertica in Eon Mode Maximizes Cloud Economics

Vertica in Eon Mode, separates the computational processes from the storage layer of the Vertica database. Running on Amazon EC2 (compute) and S3 (storage), Eon Mode demonstrates excellent performance, superior scalability, and robust operational behavior. With these innovations, Vertica 9.1 delivers on the promise of cloud economics so that data engineering teams only have to provision the compute and storage resources they need from month to month, day to day, or hour to hour, while supporting efficient elasticity as new use cases or workloads emerge.

Vertica in Eon Mode includes an intelligent, new caching mechanism on the nodes that enables the separation of compute and storage without any compromise on the speed and breadth of all Vertica analytical functions, including time series, geospatial, pattern matching, as well as the full suite of in-database machine learning capabilities to address each step in the entire predictive analytics process.

* *Techvalidate reference:* www.techvalidate.com/product-research/hp-vertica/charts/B9FBA0-073

Key Benefits

- **Cloud Economics**—Enable organizations to optimize infrastructure costs and simplify operations for their Vertica cloud deployments on AWS by separating compute resources from data storage.
- **Data Preparation for Machine Learning**—Reduce the number of features for machine learning modelling, reducing the time it takes to train a machine learning model and lowering computation cost.
- **Machine Learning Evaluation Function**—Assist Data Scientists in evaluation and choosing the right machine learning model to solve business problems with support for F-1 score, precision recall curve, and area under the curve matrix.
- **End to End Protection**—Secure Vertica data at rest, in motion, and in use via tighter integration between Vertica and Voltage SecureData.

Advanced, In-Database Analytics and Applied Machine Learning

Organizations are applying predictive analytics to everything from improving machine uptime to reducing customer churn. Vertica 9.1 provides new powerful functions for dimensionality reduction with feature extraction and additional evaluation metrics for machine learning models.

■ Singular Value Decomposition (SVD) and Principal Component Analysis (PCA)—

Popular transformation functions enable data scientists to drastically reduce the number of features for machine learning modelling, giving users the option to select threshold for cumulative explained variance. This can greatly reduce the time it takes to train a machine learning model, thereby reducing computation costs and avoiding overfitting in model training. PCA also makes it easier to visualize the data.

■ Precision Recall and F-1 Score for evaluation of machine learning algorithms—

F1 offers a single metric for the precision-recall curve to compare machine learning models.

■ Receiver Operating Characteristic function—

Vertica now supports Area under the curve (AUC) evaluation metric, the most commonly used methods for evaluating performance of binary classifiers and are particularly useful for imbalanced classes.

Tighter Integration with Micro Focus Secure Data

Vertica 9.1 includes a high-performance User-Defined Extension (UDx) that integrates Vertica and Micro Focus® SecureData to protect Vertica data at rest, in motion, and in use, by utilizing Format Preserving Encryption.

Organizations can now protect sensitive data as soon as it is acquired and ensure that the data is always used, transferred, and stored in protected form.

Performance and Management Enhancements

In addition to the performance gains and cost reductions that organizations will experience by running dynamic workloads on AWS with Vertica in Eon Mode, Vertica 9.1 includes core architectural enhancements that will improve performance of subqueries with data joins. The Vertica Management Console also includes new management advancements that enable administrators to visually monitor how their users are accessing data in HDFS data lakes by specific table, formats, files, and also measures overall data consumption capacity. This dashboard enables administrators to target frequent and business-valuable queries that could benefit from faster performance by importing HDFS data into Vertica.

Try Vertica Today

Vertica is the core SQL database analytics engine that was purpose-built with speed, scalability, simplicity, and openness. With Vertica, your queries can run 50-1,000x faster than any data warehouse or database technology.¹ It's proven to run at Exabyte-scale and gives you complete openness to use any BI/ETL tool, run as SQL on Hadoop, and leverage scalable predictive analytics and a comprehensive library of built-in advanced analytical functions.

Get started today and download Vertica Community Edition, a free version of the Vertica Advanced Analytics Platform. Store up to 1 TB of data and deploy Vertica on a 3-node cluster. Sign up for Vertica Community Edition at: www.vertica.com/try

Contact us at:
www.vertica.com

System Requirements

Server Platforms

Vertica 9.1 is supported on the following server platforms:

- Red Hat Enterprise Linux 6.6, 6.7, 6.8, 6.9, 7.0, 7.3, 7.4
- SUSE Linux Enterprise Server 11 SP3, 12 SP2, 12 SP3
- Oracle Enterprise Linux (Redhat Kernel) 6.7, 6.8, 6.9, 7.3, 7.4
- Debian Linux 7.6, 7.7, 8.5, 8.9
- CentOS 6.6, 6.7, 6.8, 6.9, 7.0, 7.3, 7.4
- Ubuntu 12.04, 14.04 LTS
- Amazon Linux 2.0 on AMI

Client Platforms

The following client interfaces and platforms are supported:

- JDBC on all Java 5, 6, 7, 8, or 9-compliant platforms
- ODBC for Windows, Linux, AIX, HP-UX, OS X
- Perl and Python on Linux, AIX, HP-UX, OS X
- ADO.NET provider for Microsoft .NET Framework 3.5 SP1 or higher
- Connectors for HDFS and HCatalog
- Connectors for Informatica and Microsoft BI Suite available from Vertica.
- Connectors for popular data integration tools including IBM, Pentaho, Syncsort and Talend from vendor. Contact vendor for details.
- Vertica ODBC driver is ODBC 3.5 compliant.
- Vertica JDBC driver is JDBC 4.0 compliant.