

Vertica 9

Vertica 9 introduces unified advanced analytics database features advancements in in-database machine learning, direct querying of Parquet data on AWS S3, support for Google Cloud Platform and Microsoft Power BI, and beta release of cloud optimized separation of compute and storage

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 provides a new way to compare how algorithms perform on any given set of data as well as a comprehensive set of new machine learning algorithms for clustering and prediction.

- Added cross-validation function, enabling data scientists to save time by comparing the performance and accuracy of the machine learning models and selecting the right one for the data set.
- Supported new data-preparation and exploration functions such as converting categorical to numerical data and statistical summary function for numeric columns, deriving greater insight from the data, while improving the quality of analysis.
- Streamlined end-to-end workflow simplifies production deployment of machine learning models—particularly for customers that embed Vertica and require the ability to replicate models across clusters.

Improved Core Data Management and High Analytical Performance

Vertica is proven to run mission-critical big data analytical initiatives at extreme scale. This latest version provides greater support of the information lifecycle with the addition of hierarchical partition management that improves query

performance at Exabyte scale. Vertica 9 also includes consistently faster refresh of live aggregate projections, improving Vertica cluster performance and speeding up node recovery.

Multi-Cloud Deployment with a Consistently Optimal Experience

As more and more customers deploy big data in the cloud, Vertica is now proven and optimized to run on Google Cloud Platform, in addition to already supported AWS and Azure. This release also makes it easier to provision and deploy Vertica nodes in AWS via the Management Console, enabling customers to easily, get started with analytics in the cloud. Further, Vertica 9 enables analytics for cloud data-lake, without having to move and re-format data via direct querying of Parquet data on S3.

Hadoop Integration for Secure Access

Furthering our ongoing commitment to open innovation and secure access, Vertica 9 introduces support for Kerberos realms that enables granular control over different business units, accessing data residing in a Hadoop data lake. This release also centralizes security policy for Cloudera and Vertica with support for Apache Sentry. This support enables security policies and privileges associated with Hadoop user to govern access control in Vertica, reducing operational burden for organizations deploying Vertica in a Hadoop environment.

Vertica 9 at a Glance

- Vertica 9 provides a comprehensive set of new machine learning algorithms for categorization, avoiding overfitting and prediction to enhance processing speed.
- Vertica introduces support for Google Cloud Platform, available via the Google Marketplace, giving organizations the flexibility and freedom to choose yet another leading cloud platform for their needs.
- Vertica 9 delivers analytics for cloud data-lake, via direct querying of Parquet data on AWS S3.
- Vertica integrates with Hadoop Sentry and Kerberos Realms for ease of management and secure access of Hadoop data-lake.
- Vertica 9 features a new Eon Mode Beta installation method on AWS. Users can test this new mode that separates compute from storage and provides rapid elastic scaling up and down of the Vertica cluster.

Vertica Eon Mode Beta for Cloud Economics

With Vertica 9, users can take part in our open beta test of a new revolutionary feature of Vertica called Eon Mode Beta. When installed in this mode, Vertica users deploy on AWS in a new way. This mode can capitalize on cloud economics through rapid compute scaling, combined with affordable S3 storage while enjoying the same fast query processing, you have come to expect from Vertica. This new architecture provides separation of compute and storage, resulting in rapid elastic scaling of

the Vertica cluster which is ideal for just-in-time workload based provisioning.

More Vertica 9 Enhancements

Vertica 9 also enables organizations to apply advanced analytics with ease of management and enterprise-class performance and reliability. Additional Vertica 9 features include:

- **Flattened Tables**—The flattened tables feature makes the task of performing complex JOINS across multiple tables much less cumbersome and much more performant. Analysts can write straight-forward, fast-running queries as if the data resided in a single big flat table, simplifying and speeding the process of big data analytics in databases with complex schemas.
- **Parquet Writer**—Introduces a new HDFS Parquet writer—built on Vertica’s fast and reliable ability to not only read, but now write data and results to HDFS—to derive and contribute immediate insights on growing data lakes in an organizations’ Hadoop data pipeline.
- **Universally Unique Identifier (UUID) as a new data type**—Allows users to store UUID columns in a space-efficient manner than having to store them as text strings.
- **Eco-system integration**—Vertica 9 natively integrates with key ecosystem technologies and open source innovation, including Microsoft PowerBI, Cloudera Manager and Apache Spark 2.1.

Vertica 9 is designed to accelerate query performance at Exabyte scale, to analyze data in the right place via direct query of Parquet data from S3, with freedom from underlying infrastructure via added support for Google Cloud Platform, to extend in-database machine learning with a comprehensive set of new machine learning algorithms for categorization, overfitting and prediction, and to support beta release of flexible cloud optimized separation of compute and storage.

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 Vertica Enterprise Edition. 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

Server Platforms

Server Platforms Vertica 9 is supported on the following server platforms:

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

Client Platforms

The following client interfaces and platforms are supported:

- JDBC on all Java 5, 6, or 7-compliant platforms
- ODBC for Windows, Linux, Solaris, AIX, HP-UX, OS X
- Perl and Python on Linux, Solaris, AIX, HP-UX, OS X
- ADO.NET provider for Microsoft .NET Framework 3.5 SP1 or higher
- Connectors for Hadoop MapReduce, HDFS, HCatalog and Pig
- 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.

www.vertica.com



Vertica Headquarters

150 Cambridgepark Drive
Cambridge, MA 02140

Learn more at: www.vertica.com