

Dig into Big Data analytics with HPE Vertica 8

New improvements to the HPE Vertica SQL analytics platform include Microsoft® Azure certification, optimized Apache Spark and Hadoop integration, and expanded in-database machine learning and analytics.

Frontloader at-a-glance

- HPE Vertica extends in-database analytics with machine learning algorithms such as k-means, linear regression, logistic regression, support vector machine (SVM), random forest and Naive Bayes that are built into the core of Vertica.
- HPE Vertica 8 supports Microsoft Azure and enables copying of Vertica data between multiple clouds for disaster recovery and development and test use cases.
- HPE Vertica features Apache Spark Connector for bidirectional data transfer for simplified integration with real-time architectures
- Seamless backup and restore of Vertica data to and from AWS S3
- HPE Vertica offers high-performance support for Hadoop file formats such as ORC and Parquet. Native reader for ORC and along with a reader and writer for Parquet.
- HPE Vertica 8 features many more enhancements to make managing and securing your data for analytics simple

High performance analytics at Exabyte scale with complete reliability

HPE Vertica runs mission-critical Big Data analytical initiatives at extreme scale. In data-driven organizations, thousands of concurrent users access it and extract meaningful insights. Hewlett Packard Enterprise announces the release of HPE Vertica 8, codenamed “Frontloader,” enhancing its open and scalable capabilities.

This latest version features advanced, in-database analytics, optimized open source integration, additional options for cloud deployment, and much more.

Advanced, in-database analytics, and applied machine learning

Organizations are applying predictive analytics to everything from improving machine uptime to reducing customer churn. Today, data scientists in these organizations are challenged with processing large datasets due to memory or computational limitations of R and Spark, and moving these large datasets between systems to operationalize models. As a result, data scientists are building and tuning models based on a subset of data—not the true representative dataset—and experiencing delays in operationalizing and deploying the predictive models into production.

With Frontloader, HPE Vertica introduces support for native machine learning

algorithms. Data scientists can now leverage SQL to create and deploy machine learning models natively based on larger datasets without downsampling to accelerate the decision-making process.

Optimized and integrated Spark, Kafka, and Hadoop innovations

Furthering our ongoing commitment to open innovation, HPE Vertica 8 introduces Apache Spark Connector for bidirectional data movement, improves monitoring of Apache Kafka topics, and accelerates SQL on Apache Hadoop support. The Spark Connector exploits HPE Vertica’s parallelism providing fast, scalable data transfer between Spark and HPE Vertica, benefitting Spark users to leverage HPE Vertica’s advanced SQL analytics. For Hadoop deployment, users now have faster access to HDFS data for running SQL queries and a more secure access mechanism via Kerberos.

Further, HPE Vertica now boosts the performance of queries that run on Apache Parquet and ORC files. This latest release of HPE Vertica embraces and extends open source technologies with native file format support, enabling organizations to harness their entire data to improve business outcomes.

Multi-cloud deployment with a consistently optimal experience

As more and more customers are deploying Big Data in the cloud, HPE Vertica is now

Solution brief

Server platforms

Frontloader is supported on the following server platforms:

- Red Hat Enterprise Linux® 6, 7
- SUSE Linux Enterprise Server 11
- Oracle Enterprise Linux (Red Hat Kernel) 6, 6.8
- Debian Linux 7.5 up to and including 7.7
- CentOS 6, 7
- 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 HPE Vertica
- Connectors for popular data integration tools including IBM, Pentaho, Syncsort, and Talend from vendor; contact vendor for details.
- HPE Vertica ODBC driver is ODBC 3.5 compliant
- HPE Vertica JDBC driver is JDBC 4.0 compliant

Check the release notes for updates to the supported configurations.

[TechValidate survey of 219 users of HPE Vertica, December 2015.](#)

Learn more at
hpe.com/vertica



Sign up for updates

**Hewlett Packard
Enterprise**

proven and optimized to run on Microsoft Azure cloud, in addition to already supported Amazon Web Services (AWS). This release also offers expanded AWS support with seamless access to Amazon S3 object store, backup and restore to and from S3, integration with AWS CloudWatch for monitoring and tighter security. With the flexibility to start small and grow as your business grows, HPE Vertica now enables you to transition your data warehouse to the cloud, between clouds, on-premises, and back seamlessly.

More Frontloader enhancements

HPE Vertica 8 also enables organizations to apply advanced analytics with enhanced security and enterprise-class performance and reliability. Additional HPE Vertica 8 features include:

- **Advanced system monitoring**—enables organizations to perform historical analysis of log data and faster rendering of health status in the management console by continuously streaming, monitoring information from production cluster to an external data mart on a standalone cluster.
- **Geospatial analytics enhancements**—HPE Vertica 8 offers users the ability to perform geospatial joins without the need to transform data to Vertica's unique geometry and geography data types, accelerating data analysis.
- **Fast lightweight table copy for Extract Load Transform (ETL)**—the release also provides administrators the ability to create writeable, space-efficient linked copy for a given source table, significantly speeding up the copy process.
- **FIPS compliance**—the release conforms with FIPS 140-2 (Level 1) standard for secure communication between the user and HPE Vertica nodes by dynamically using FIPS-certified, OpenSSL cryptographic module provided by the Red Hat® operating system.
- **Python user-defined scalar function (UDSF)**—users can now develop UDSFs in

the Python language using the HPE Vertica Python SDK, extending the choice of languages for deploying predictive analytics at speed and scale.

- **Access control enhancement**—improves connection usability when there is a large group of users by defining connection limits at a user level, denying connection when the limit is exceeded and closing any open sessions that have been idle for too long.

HPE Vertica 8 is designed to speed the loading of tables with hundreds of columns by distributing the load among nodes, resulting in significant performance improvement. Our latest release is resilient for demanding environments. HPE Vertica can execute queries with a modified query plan even when a node fails—ensuring queries run fast, no matter what challenges occur.

HPE Vertica leads the way in Big Data analytics

HPE Vertica is at the core of SQL database analytics engine of the HPE Big Data Platform that was purpose-built with speed, scalability, simplicity, and openness. With HPE 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 business intelligence or ETL tool, run as SQL on Hadoop, and leverage scalable predictive analytics and a comprehensive library of built-in analytical functions.

The software delivers comprehensive advanced SQL analytics. As a standards-based relational database, it includes full support for SQL, JDBC, and ODBC. This allows users to preserve years of investment and training in these technologies because SQL programming tools and languages work seamlessly together. HPE Vertica also includes a wide array of built-in analytical functions, including geospatial, time series, pattern matching, and more.

© Copyright 2017 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Apache Hadoop, Hadoop, and Apache are either registered trademarks or trademarks of the Apache Software Foundation in the United States and/or other countries. All other third-party trademark(s) is/are property of their respective owner(s).

a00000387ENW, July 2017, Rev. 1