



**Hewlett Packard
Enterprise**

HPE Vertica Enterprise

Big Data analytics with no limits and no compromises

High-performance, Big Data analytics

Big Data analytics solutions shouldn't pose a big obstacle to decision making. With HPE Vertica Enterprise, you get high-performance SQL analytics for all of your data to make better business decisions now, not after your opportunity has passed.

Packed with powerful features for harnessing the value from all your data, HPE Vertica Enterprise manages massive amounts of data quickly and reliably, without compromise. As part of the same trusted core of the HPE Vertica Analytics Platform, HPE Vertica Enterprise enables you to perform queries 50 to 1,000 times faster than traditional databases, at a fraction of the cost, using a fraction of the hardware.¹ Unlike traditional RDBMS databases, which are not designed for data analytics and today's complex analytics workloads, the HPE Vertica Analytics Platform is built from the ground up to deliver breakneck performance on petabyte volumes at nearly infinite scale.

Key features and benefits

Columnar storage and execution

By eliminating costly disk I/O associated with traditional row-oriented SQL databases, you can perform queries 50 to 1,000 times faster.¹ Store 10x–30x more data per server than traditional databases with patented columnar compression.²

“Scale-out” massively parallel processing (MPP)

Scale your Big Data analytics solution by adding industry-standard servers and storage to your analytics environment.

Extensible in-database analytics framework

Achieve open access to in-database processing through a robust development framework for procedural, user-defined analytics. In addition to using built-in SQL analytic and aggregate functions, you can define your own custom data analytics functions with a software developer's kit (SDK). The SDK features secure sandboxing, and resulting functions run in parallel for fast performance.

In-database analytics library

Out-of-the-box, in-database analytics include event-series pattern matching, event-series joins, linear regression, geospatial support, advanced time series, and many more. You get open source analytics libraries, including thousands of packages from Comprehensive R Archive Network (CRAN).

Native support for Hadoop, Kafka, and more

HPE Vertica Enterprise includes application programming interfaces (APIs) for user-defined aggregates, analytics, and multi-phase transform functions. These dynamically integrate with Apache Hadoop, Apache Kafka, and MapReduce, so that you can analyze large sets of structured, semi-structured, and unstructured data fast.

Automatic high availability

Run non-stop with data replication, failover, and recovery; HPE Vertica Enterprise is

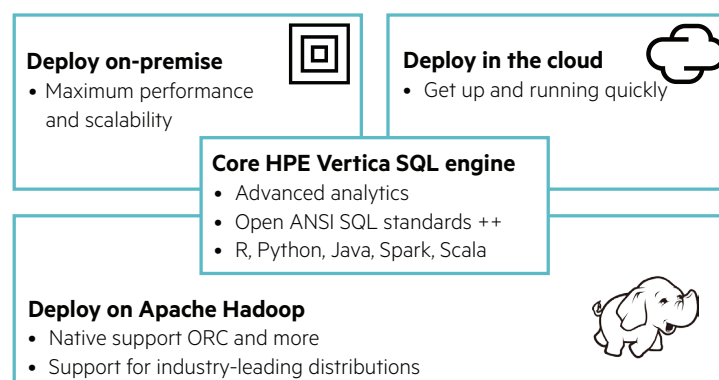
optimized for performance and is transparent to your integration and operations teams.

Optimized system and performance management

Experience a robust set of APIs for monitoring your system's resources, background processes, workload, and performance. The Management Console displays tuning recommendations from the Workload Analyzer tool, addressing a wide range of system concerns, including empty administrator passwords, high CPU usage, and recommendations for better data management.

Broad deployment and consumption models—all based on the same Vertica core

Available on-premise, in the cloud, or on Hadoop, HPE Vertica offers proven Big Data analytics that can deliver unmatched speed and scale. Depending on your use case, if you need extra capacity and have no time to stand up on-premise hardware, deploying Vertica in the cloud is an attractive option. For organizations that prefer to apply a sophisticated SQL query engine to their growing data lakes, HPE Vertica for SQL on Apache Hadoop is another popular Vertica consumption model.



¹ TechValidate reference techvalidate.com/product-research/hp-vertica/charts/B9F-BA0-073

² An HPE internal study

Out-of-the-box productivity

HPE Vertica Enterprise was designed to get your users up and running quickly, requiring fewer resources than traditional Big Data analytics solutions.

With real-time loading and querying, automated database design, and familiar interfaces, your users are productive immediately.

Operating on industry-standard hardware and with an optimized architecture, HPE Vertica Enterprise can deliver significant savings while providing you with industry-leading analytics capabilities and performance.

New technology, not retrofitted solutions

Designed and developed in this century, HPE Vertica Enterprise powers the largest organizations and most innovative business models around the world.

Real-time Big Data analytics

Gain insights into your data in real time. Consume, analyze, and make informed decisions at the speed of business.

Fast time to value

Monetize your data in a matter of minutes, not weeks or months. Prevent data blind spots from increasing your risk and uncertainty around regulatory compliance.

Agile deployment

Get increased flexibility—deploy HPE Vertica Enterprise on industry-standard hardware, virtual machines, on Hadoop, or in the cloud.

You can get all of these benefits with a leading data analytics solution that scales infinitely without compromising performance.

HPE Vertica Enterprise offerings

The HPE Vertica Enterprise comes in two editions. The Express edition consists of the base functionality and the Premium edition has additional advanced capabilities as shown below.

Enterprise capabilities	Express (Price per TB)	Premium (Price per TB)
MPP architecture for infinite scalability	✓	✓
High Availability provides tolerance for node failures	✓	✓
Role-based security	✓	✓
User-Defined Extension (UDx) for filtering, parsing, scalar functions, aggregates, transforms, and more	✓	✓
Standard SQL (ANSI 99)	✓	✓
Flex Tables—load semi structured data (JSON, Avro, etc.) into a table for querying with pre-built parsers	✓	✓
Workload Analyzer, DB Designer, and Management Console for ease of management and optimal query performance	✓	✓
Elastic Cluster for seamlessly adding or removing nodes in a cluster	✓	✓
Advanced SQL analytics (Time Series, SQL Windowing, Gap Filling, more)* and Text Search		✓
Routable Query (Key Value) Interface to achieve high throughput for low-latency and high-concurrency queries		✓
Fault Groups provides high availability and reduces risk of correlated failures		✓
Advanced Analytics with Sentiment, Geospatial, R Extensions, and In-database Machine learning		✓
Advanced Relational OLAP functions (Rollup Aggregate, Grouping Sets Aggregate, Cube Aggregate, and Pivot)		✓
Live Aggregate and Pre-Join Projections to optimize query performance for frequently used queries containing aggregates		✓

Technical specifications

Supported operating systems	64-bit operating systems on x86, x64 architecture: <ul style="list-style-type: none"> • Red Hat® Enterprise Linux® 6.5 up to 6.7 and 7.0 • SUSE Linux Enterprise Server 11.0 SP3 • Oracle Enterprise Linux 6 (Red Hat-compatible kernel only) • Debian Linux 7.5 and up to 7.7 • CentOS 6.5 up to 6.7 and 7.0 		
Supported file systems	Extensively tested on all supported Linux platforms running ext3 or ext4 file systems. ³		
	Operating system	Processor	Driver manager
ODBC drivers (32-bit and 64-bit drivers provided)	Microsoft® Windows® <ul style="list-style-type: none"> • Windows 7, 8, and 10 all variants • Windows 2008 R2 all variants • Windows 2012 	x86, x64	Microsoft ODBC MDAC 2.8 iODBC 3.52.6 or later unixODBC 2.3.0 or later DataDirect 5.3 and 6.1 or later
	Microsoft Windows 2012	x86, x64, IA64	
	Red Hat Enterprise Linux 6 and 7	x86, x64	
	SUSE Linux Enterprise 11	x86, x64	
	Oracle Enterprise Linux 6 x64 (Red Hat-compatible kernel only)	x86, x64	
	Ubuntu 12.04 and 14.04 LTS	x86, x64, IA64	Microsoft ODBC MDAC 2.8 iODBC 3.52.6 or later unixODBC 2.3.0 or later DataDirect 5.3 and 6.1 or later
	CentOS 6 and 7	x86, x64	Microsoft ODBC MDAC 2.8 iODBC 3.52.6 or later
	Solaris 10	x86, x64, SPARC	unixODBC 2.3.0 or later DataDirect 5.3 and 6.1
	AIX 5.3, 6.1	PowerPC	
	HP-UX 11i V3	IA32, IA64	
	Mac OS X 10.9	x86, x64	Mac OS X Driver Manager (iODBC)
JDBC drivers (JDBC 3.0 compliant)	Supported on all operating environment platforms		
	Operating system	Processor	.NET requirements
ADO.NET driver	Microsoft Windows <ul style="list-style-type: none"> • Windows 2008 R2 all variants • Windows 2012 • Windows 7, 8, and 10 all variants 	x86, x64	Microsoft .NET—Framework 3.5 SP1 or later
vsq client	Supported on all HPE Vertica Enterprise Edition ODBC driver platforms		
	Perl version	Perl driver modules	ODBC requirements
Perl	5.8 5.10	DBI driver version 1.609 DBD: ODBC version 1.22	See ODBC drivers.
	Python version	Python driver modules	ODBC requirements
Python	2.4.6 2.7.3 3.3.4	pyodbc 2.1.6 pyodbc 3.0.6 pyodbc 3.0.7	See ODBC drivers.

Data sheet

Technical specifications (continued)

R version			
Plug-in version	Operating system	PowerCenter version	
R	3.0.0 or higher		
Plug-in for Informatica PowerCenter	Plug-in version	Operating system	PowerCenter version
	7.x	Microsoft Windows <ul style="list-style-type: none">• Windows 2003 and 2003 R2 all variants• Windows 2008 and 2008 R2 all variants• Windows 7 all variants Red Hat Enterprise Linux 5 (32- and 64-bit), Solaris, AIX, HP-UX	9.x
Hadoop Connector	Supports the following combinations of Apache Hadoop and Apache Pig: <ul style="list-style-type: none">• Hadoop 2.0.0 and Pig 0.10.0 Supports Cloudera Distribution Versions: <ul style="list-style-type: none">• Cloudera Distribution Including Apache Hadoop (CDH) 4		
Hadoop distributions	Cloudera (CDH) 5.5, 5.6, 5.7 Hortonworks Data Platform (HDP) 2.2, 2.3, 2.4 MapR 4.1, 5.0		

³ While other file systems have been successfully deployed by some customers, Hewlett Packard Enterprise Company, (HPE) cannot guarantee performance or stability of the product on these file systems. In certain support situations, you may need to migrate from these untested file systems to help you troubleshoot or fix an issue. HPE Vertica Community Edition does not support Linux Logical Volume Manager.

Try HPE Vertica today

Get started today and download HPE Vertica Community Edition, a free version of HPE Vertica Enterprise Edition. Store up to 1 TB of data and deploy HPE Vertica on a three node cluster. Sign up for HPE Vertica Community Edition at my.vertica.com.

Learn more at
hpe.com/software/vertica



© Copyright 2012, 2016 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.

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. Oracle is a registered trademark of Oracle and/or its affiliates. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Apache Hadoop and Hadoop are either registered trademarks or trademarks of the Apache Software Foundation in the United States and/or other countries.

4AA0-6143ENW, August 2016, Rev. 1