



HP Vertica Analytics Platform 7

Shine a Light on Dark Data



HP Vertica Analytics Platform 7

HP Vertica Analytics Platform 7 “Crane”

The next-generation analytics platform that shines a light on dark data:

- Semi-Structured Data Exploration and Visualization
- The Richest and Most Open SQL-on-Hadoop and Expanded Data Management
- Simplified Cloud Deployment
- Enhanced End-to-End Security
- New Java Extensibility
- Much More

Organizations recognize the immense potential value in big data. However, it can be particularly challenging to find patterns in or explore new forms of semi-structured data such as social media, web logs, sensors and more. These newer, yet rapidly growing and critically important, data types require a lengthy process to load into traditional analytic platforms and data warehouses before they can be explored, analyzed, and ultimately monetized -- remaining “dark” to many organizations.

The HP Vertica Analytics Platform 7 “Crane,” the latest major release of the industry’s leading analytics platform that is purpose-built for big data, enables you to accelerate business value from a vastly expanded variety of data by simplifying the exploration and analysis of semi-structured and “dark data.” This release also includes the richest, most open SQL on Hadoop, critical security and performance enhancements, and much more.

This document highlights key features of the release. For a complete list of features and enhancements, see the HP Vertica 7 Enterprise Edition New Features Guide.

Sign up for a free evaluation

my.vertica.com/evaluate/



Share with colleagues



Semi-Structured Data Exploration and Visualization

HP Vertica Flex Zone

HP Vertica 7 lets you shine a light on dark data simply and cost-effectively with HP Vertica Flex Zone, a recommended and revolutionary new offering. HP Vertica Flex Zone gives you the power to quickly and easily load, explore, analyze, and monetize emerging and rapidly growing forms of semi-structured data such as social media, sensor, log files, and machine data. Use your favorite industry-standard business intelligence (BI) and visualization tools for exploration of Flex Zone data without creating schemas. HP Vertica Flex Zone offers patent-pending auto-schematization™, eliminating the need for coding-intensive schemas to be defined or applied before the data is loaded for exploration. This one-step schema capability enables schemas to be created and applied as-needed for high-performance analytics and lets you address the ever-changing structure of data with ease.

An open approach allows users to deploy many types of high-performance dashboards to visually explore and draw conclusions from data patterns across a full spectrum of structured and semi-structured data. Analysts, data scientists, and business users can now explore and visualize information such as JSON and text without burdening or waiting for their IT organizations to extract, structure, and load the data.

The Richest and Most Open SQL on Hadoop and Expanded Data Management

The Richest and Most Open SQL on Hadoop

Data stored in the Hadoop Distributed File System (HDFS) offers the potential for extraordinary business insight yet can be challenging to explore in traditional database platforms. HP Vertica 7 delivers the industry's most open "SQL-on-Hadoop" capabilities. Unlike other SQL-on-Hadoop platforms, HP Vertica works with major Hadoop distributions, ensuring high-performance analytics across the broadest range of data types and sources.

Leverage HCatalog with Vertica

HP Vertica 7 supports HCatalog, Hadoop's table and storage management layer. This integration enhances your ability to navigate data in Hadoop and quickly and easily load or query that data from the HP Vertica Analytics Platform. Specifically, this feature lets you access your data without worrying about where the data is stored, or changing schemas in your big data files. You no longer need to spend time recreating schema definitions in HP Vertica when the data changes. In addition, you can more easily explore external tables stored in Hadoop.

Simplified DBA Operations

To make life easier for DBAs, HP Vertica 7 includes enhancements in the installation and management process. This release also includes simplified administration with log management updates. The DBA now can specify how long to keep history to better deal with compliance requirements on data retention.

Enhanced Data Loading

With HP Vertica 7, you can load data with speed and ease. Data managers can now use multiple machines to parse data while loading, improving performance of loading massive amounts of data into the HP Vertica Analytics Platform. The Parallelized Loader makes loading very large data files easy by dividing file parsing tasks across CPUs, resulting in improved load speeds.

While loading data, a complete data set ensures accurate and reliable results. With improved loading, you can now query on records that the loading process failed to load. You can use this table and address data anomalies or exceptions during the load. Finally, Windows users can achieve better control and performance on ODBC drivers by changing advanced ODBC options in the Windows ODBC Administrator without changing the registry.



HP Vertica Analytics Platform 7 "Crane"

Shine a light on dark data

Enhanced Management Console

The HP Vertica Management Console is a web-based work environment for managing your HP Vertica cluster. Using only a web browser, you can create, import, manage, and monitor multiple databases on one or more clusters. You can manage users and assign access to the cluster. Get a better understanding of any technical issues by browsing through alerts and metrics, and then use the Management Console to dynamically tune your database for the fastest analytics. The HP Vertica Management Console is designed to make it easy to manage the HP Vertica Analytics Platform.

Workload Analyzer

HP Vertica 7 enhances the Workload Analyzer (WLA), which analyzes system information held in SQL system tables. WLA intelligently monitors the performance of your SQL queries, workload history, resources, and configurations to identify the root causes of poor query performance. It then returns a set of recommendations that let you quickly and easily tune query performance without needing sophisticated skills. New features give more precise CPU control and better performance for certain types of analytical requests.

Simplified Cloud Deployment

New Amazon Machine Image (AMI)

To improve support for Amazon EC2 deployments, HP Vertica 7 provides a Hardware Virtual Machine (HVM) based Amazon Machine Image (AMI). Using Amazon's Cluster Compute instance, you can leverage all the benefits of cloud computing with the rich analytics functionality, scalability, and performance benefits of HP Vertica. The AMI includes a new set of Cloud Scripts that simplify cluster installation, configuration, and deployment. The Cloud Scripts automate a range of activities, such as creating clusters from pre-loaded machine images, adding or removing cluster nodes, and changing machine IP addresses—all of which enable more agile cluster management.

Enhanced End-to-End Security

Secure Data Access

To provide you with the highest level of data security and compliance, HP Vertica 7 offers secure access to your data stored in Vertica and HDFS with Kerberos authentication. This secure HDFS Connector ensures security by mapping HP usernames into Kerberos principals for authentication against the Kerberos server. To validate access, Hadoop is presented with Kerberos authentication information.

Secure Client Drivers

HP Vertica 7 also offers Kerberos authentication to HDFS, HCatalog, ODBC, JDBC, ADO.NET, and Windows VSQLE drivers via the GSS-API. Kerberized client drivers enable secure access from HP Vertica clients to HP Vertica servers via Kerberos, ensuring the highest level of security compliance for your data.

Get connected

hp.com/go/getconnected

Current HP driver, support, and security alerts delivered directly to your desktop

© Copyright 2013 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP 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. HP shall not be liable for technical or editorial errors or omissions contained herein.

Trademark acknowledgments, if needed.

4AA4-xxxxENW, Created Month 2012



New Java Extensibility

Java SDK

Java remains one of the most popular languages for developing extensions. With the new Java SDK, you can now create in-database extensions in Java that harnesses the performance and scalability of HP Vertica's columnar, MPP architecture. In addition, data access using the JDBC 4.0 API is now approved for use in HP Vertica 7. Finally, Java developers can take advantage of a new Java User-defined Loader (UDL) API, which allows them to define and parse files for use within HP Vertica.

And much more...

In addition to the above features, this release also provides additional platform enhancements to help you gain the most out of your investment. For a full list of features, enhancements, and changes please see the HP Vertica 7 Release Notes and New Features Guide.

Server Platforms

HP Vertica 7 is supported on the following server platforms:

-
- Red Hat Enterprise Linux 5, 6
 - SUSE Linux Enterprise Server 10, 11
 - Oracle Enterprise Linux (Redhat Kernel) 6
 - Debian Linux 5, 6
 - CentOS 5, 6
 - Ubuntu 12.04 LTS
-

Resources

For data sheets, case studies, white papers, videos, and recorded Webinars
www.vertica.com/resources

See more product details at
www.vertica.com/the-analytics-platform

Region

HP
150 Cambridgepark Drive
Cambridge, MA 02140
sales@vertica.com
617-386-4400

Client Platforms

The following client interfaces and platforms are supported:

-
- JDBC on all Java 5-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 or higher
 - Connectors for Hadoop MapReduce, HDFS, HCatalog and Pig
 - Connectors for Informatica and Microsoft BI Suite available from HP Vertica.
 - Connectors for popular data integration tools including IBM, Pentaho, Syncsort and Talend from vendor. Contact vendor for details.

HP Vertica ODBC driver is ODBC 3.5 compliant.

HP Vertica JDBC driver is JDBC 4.0 compliant.

Learn more at

www.vertica.com/crane