

HP Vertica Analytics Platform 6.1



The Next-Generation Analytics Platform Purpose Built to Optimize Big Data

HP Vertica Analytics Platform 6.1

The HP Vertica Analytics Platform 6.1 is the latest update to the real-time analytics database platform purpose built for Big Data. This release includes features and enhancements focused on delivering high-performance analytics, tighter integration with Hadoop, expanded data management, simplified deployment to the cloud, and enhanced platform usability and administration. As a result, you can now get even faster answers from your Big Data with improved support for Hadoop and R, a deployment model that fits your business need, and complete scalability and low cost as your data grows.

This document highlights key features of the release. For a complete list of features and enhancements, see the HP Vertica 6.1 Enterprise Edition New Features Guide.

Enhanced Support for User-Defined Functions in R

With the HP Vertica 6 release, you can deploy functions written in R on the HP Vertica Analytics Platform. HP Vertica 6.1 expands this functionality with a pre-built R language pack to help you shorten the time to become productive with R. Use new support for parameterized and polymorphic R functions to write generic and flexible functions that, through support of a variable number and types of arguments, can handle a variety of data sets. Improve the performance of your R functions with new features to specify volatility and null-input settings.

New and Improved SQL Functions

HP has enhanced the HP Vertica Analytics Platform through expanded and improved performance of SQL-99 functions. This release includes support for the SQL-99 WITH clause as well as set operators such as INTERSECT, MINUS, and EXCEPT. These new operators simplify users' SQL expressions and ease porting of SQL scripts from other database products.

Hadoop Distributed File System (HDFS) Connector

HP Vertica 6.1 extends the platform's Hadoop integration capabilities with a connector for loading data directly from Hadoop Distributed File System (HDFS), affording you a "best-of-both world's" approach to monetizing your Big Data. With this new connector, you can load data from HDFS using the HP Vertica native COPY facility. This mechanism simplifies and accelerates the process of loading data stored in HDFS without any MapReduce coding. The connector also ensures that data is loaded from the Hadoop cluster with the optimal amount of parallelism. By using the connector with the HP Vertica External Tables feature, you can even query data in HDFS without copying it into HP Vertica.

Tiered Storage

HP Vertica 6.1 now provides more control over where the platform stores data. By extending HP Vertica Storage Location capabilities, you can now label storage tier locations and define policies for storing data objects on these tiers. These features enable flexible storage tiering strategies that make optimal use of your enterprise's storage media. This feature is especially useful for deployments that have access to a range of storage devices such as FLASH and HDDs and SANs.

Data Partition Archiving

Building on the platform's rich partitioning features, HP Vertica 6.1 enables the archiving of partitions to offline storage. Unlike data backup, partition archiving moves selected partitions out of HP Vertica online storage rather than creating a backup copy of the data. This capability simplifies the archiving process for historic data not needed in online storage. Archived data can then be restored to online storage as needed.

Simplified Cloud Deployment

To improve support for Amazon EC2 deployments, HP Vertica 6.1 now provides a Hardware Virtual Machine (HVM) based Amazon Machine Image (AMI). Using Amazon's Cluster Compute instance, you can leverage all of 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.

Deployment Validation Tools

To assist you with validating deployment environments, HP Vertica 6.1 introduces new utilities that validate cluster node configurations. These tools validate that the nodes' I/O, CPU, and network subsystems have been configured sufficiently to support an HP Vertica deployment. The tools can be run from the HP Vertica Management Console or directly from a command line and provide detailed measurements that can assist with hardware configuration activities.

Role-Based Access Control

The HP Vertica Management Console now provides role-based access control for improved data and system security. With this new capability, administrators have more control in securing administrative functions provided by the console. As a result, administrators can ensure organizations comply with company and industry security policy regulations.

Upgrade-Friendly SDK Extensions

HP Vertica 6.1 provides improved compatibility for extensions built with the HP Vertica SDK. User-defined extensions will function on subsequent HP Vertica versions with minimal or no porting effort required. This compatibility eases the upgrade process and ensures your HP Vertica extensions will continue to function as written.

And much more...

In addition to the above features, this release also provides additional platform enhancements, including security, performance, robustness, and customer-reported issues. For a full list of features, enhancements, and changes please see the HP Vertica 6 Release Notes and New Features Guide.

Server Platforms

HP Vertica 6.1 is supported on the following server platforms:

- Red Hat Enterprise Linux 5, 6
 - SUSE Linux Enterprise Server 10, 11
 - Oracle Enterprise Linux 6
 - Debian Linux 5, 6
 - CentOS 5, 6
-

Client Platforms

The following client interfaces and platforms are supported:

- JDBC on all 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 3.0
- Plug-in for Informatica PowerCenter
- Connector for Hadoop, HDFS, and Pig

HP Vertica ODBC driver is ODBC 3.5 compliant.

HP Vertica JDBC driver is JDBC 3.0 compliant.
