

Deriving Greater Value from Your Enterprise Data Warehouse

Vertica delivers speed without compromise, scale without limits, and the broadest range of consumption and deployment models.

"We've replaced a MySQL database with Vertica. Everyone is amazed at how quickly we can get the data into the database and how fast we can query and analyze the data once loaded. What used to take days can now be done in minutes!"

IT PROFESSIONAL

Large Enterprise Health Care Company

The core, unified architecture supports all leading BI and visualization tools and works with your current ETL tools to empower your analytics. Vertica helps you derive more value from your Enterprise Data Warehouse and get to market faster with your analytics initiatives.

No Walls, No Boundaries

Many database administrators find themselves always trying to squeeze out marginal performance gains from their legacy data warehouse. In many environments, this means investing in more hardware and services. Admins will even wall-off the data warehouse in order to protect performance. This creates a boundary to access analytics for those who need it. Users expect access to reports and analytics so that they can make informed decisions.

As increases in need for concurrency and performance grow, so will the costs to license a typical data warehouse. Costs skyrocket as additional licenses, hardware and consultants need to be brought in to bring the solution up to speed.

What Is Vertica?

The Vertica Analytics Platform is purpose built from the very first line of code for the modern data warehouse. It is designed for use in data warehouses and other big data workloads where speed, scalability, simplicity, and openness are crucial to the success of analytics. Vertica relies on a tested, reliable distributed architecture and columnar compression to

deliver blazingly fast speed. A simplified license and the capability to deploy anywhere delivers on the promise of big data analytics like no other solution.

Vertica provides you with the broadest range of deployment models, so that you have complete choice as your analytical needs evolve:

- **Vertica Enterprise**—The core "shared nothing," distributed analytical database designed to work on clusters of cost-effective, off-the-shelf servers in your data center with unparalleled performance and extreme scale.
- **Vertica in the Clouds**—Optimized a pre-configured to run on AWS, Microsoft Azure, and VMware clouds, Vertica is also available as a BYOL (Bring Your Own License) model to enable you to transition your data analytical workloads to the cloud to on premise and back seamlessly.
- **Vertica for SQL on Hadoop**—Run the industry's most comprehensive Vertica SQL analytics engine directly on your Hadoop cluster and tap into advanced SQL on Hadoop capabilities, complete 100 percent of the TPC-DS queries without modification, achieve greater concurrency, and run on any Hadoop distribution.

Deliver Faster Analytics

Users don't want to wait for results. Your analytical database should provide the scalability to meet service-level agreements (SLAs) and

demanding timeframes for executing a query. With Vertica, you can gain insights into your data in near-real time by running queries 50-1,000% faster than legacy enterprise data warehouses. Because operations that took days now take hours and hours now take seconds, your analytics team can be more productive and answer business-critical questions on the spot.

Analytical Functions in Vertica

The analytical functions in Vertica span from standard SQL-99 conventions to value-added SQL capabilities to in-database machine learning. Vertica includes geospatial, time-series, event-based windows, pattern matching analytics and more. In addition, Vertica is built to be extensible. It allows the development of custom SQL analytics functions that can leverage the Vertica MPP architecture and data locality computation. Vertica supports popular languages—C++, Java, and R—to accelerate the development of user-defined extensions (UDx).

Perform In-Database Machine Learning

Vertica offers a robust set of advanced in-database analytics and machine learning functions and algorithms so that you can conduct the analytics computations closer to the data, and get immediate answers from a single place without the need to extract information to a separate environment for processing. You can use the Vertica engine to training the machine learning models without having to move data, thus avoiding the common practice of moving data, or subsets of data, from the data warehouse into another solution in order to train.

Worry Less about Schemas

For decades, it's been widely accepted that snowflake and star schemas facilitate getting optimal performance from your data warehouse. The schemas are configured based on the common analytics that your company needs. Vertica fully supports these schema types, just as any enterprise-grade database would. However, Vertica offers a flattened tables feature, where you can easily create a flat table that derives values from multiple fact and dimension tables. This simplifies the way that the analysts write their queries and boosts the performance of analytics over performing JOINS. Vertica's flattened table feature may save you from having to refactor your database as the types of analytics you run evolve. It also helps you create a multi-purpose data

warehouse without having to create operational data stores and other copies of the data.

Integrate with Existing BI, ETL Tools

If your data warehouse infrastructure relies on extract, transform, load (ETL) tools or SQL-based visualizations, your analytics platform should provide robust and powerful SQL and also be certified to work with all of your tools—not just those from your primary vendor. Every release of Vertica is certified and tested with visualization and ETL tools. This enables users to preserve years of investment and training in these technologies because all popular SQL programming tools and languages work seamlessly. Leading BI and visualization tools are tightly integrated, such as Cognos, Looker, MicroStrategy, Tableau, and others and so are all popular ETL tools like Informatica, Talend, Pentaho, and more.

Leverage Your Open Source Solutions

With an ecosystem friendly architecture that supports Apache Kafka, Apache Spark, Apache Hadoop, Python and more, Vertica brings enterprise level data analytics to your open source projects. Analysts can reach across from the data warehouse to the data lake to perform analytics and JOINS.

Real-World, Bottom Line Results

Here are some examples of organizations that have capitalized on their most strategic asset—their data—with Vertica:

Anritsu—Replaced their legacy EDW with Vertica as their new embedded analytics database to implement predictive analytics solutions that had been theoretical until now and achieved 351% ROI with a payback of just 4 months.

Guess—Delivers essential daily store reports via mobile devices for accurate sales tracking, improvements in merchandise allocation and distribution, and insightful customer purchasing behavior.

Try it and make your concept a reality. The Vertica next-generation high-performance SQL analytics engine is available as three integrated offerings to meet your varying needs—on premise, in the cloud, or on Hadoop. Your needs are unique, so your analytics database should be too. Evaluate Vertica today at: www.vertica.com/try

Contact us at:
www.vertica.com

System Compatibility

- Red Hat Enterprise Linux 6.6, 6.7, 6.8, 6.9, 7.0, 7.3
- SUSE Linux Enterprise Server 11.0 SP3*, 12 SP2
- Oracle Enterprise Linux 6.7, 6.8, 7.3 (Red Hat compatible kernel only)
- Debian Linux 7.6, 7.7, 8.5
- CentOS 6.6, 6.7, 6.8, 6.9, 7.0, 7.3
- Ubuntu 14.04 LTS

Deployment Options (Bring Your Own License)

- On-premises
- Amazon AWS Cloud
- Google Cloud
- Microsoft Azure