

MONITIZE

'Big Data' with real-time analytics

The HP Vertica Analytics System

Solution brief

One of the key aspects of the emerging Internet—where real-world objects are connected—is the massive amount of new data on the Web that will result. We live in an analytics world with more and more data to analyze arriving continuously. The sheer volume and velocity of the data means that many companies can no longer afford overnight batch load windows in order to extract information from their data.

Mobile devices—smart phones, tablets, and the like—ignited the data explosion. Social networking, gene sequencing, electronic patient records, medical imaging, added regulations, and geophysical exploration provide even more fuel.

Despite these challenges, companies have been finding there's a lot of value to the data, but the key is finding more effective and efficient ways to use real-time analytics on 'Big Data' to maximize business value.

Evolution of analytic infrastructure

'Big Data' is a term applied to data sets whose size is beyond the ability of commonly-used software tools to capture, manage, and process the data within a tolerable elapsed time.

The size, scope, and structure of that 'Big Data' is beyond the reach or capacity of traditional BI/EDW systems to store or analyze. Fortunately, the revolutionary HP Vertica Analytics System delivers a solution that can handle more data, provide more insight, and deliver more value now by harnessing the power of real-time analytics to fuel business innovation and competitive advantage.

Accelerating time-to-business value with real-time analytics

To deliver the industry's best analytics, it is essential to optimize at each layer of the analytics infrastructure—server, storage, network, and management.

By combining the powerful HP Converged Infrastructure—an architectural blueprint for the data center of the future—with the next-generation architecture of the HP Vertica Analytics Platform, HP is able to deliver ground-breaking scalability and real-time analytics performance that traditional BI/EDW systems simply cannot handle.

Built from the ground up for analytics

The HP Vertica Analytics System is the industry's most modern real-time analytics platform purpose-built on the four "Cs" to handle the volume and velocity of today's 'Big Data' workloads.

- **Columnar storage and execution**—Ideal for read-intensive (query) workloads while dramatically reducing disk I/O by orders of magnitude. It delivers 50-1000x faster query performance than traditional row-store databases. Vertica stores data the way it's queried for best performance.
- **Capacity optimization**—Providing more compression options when similar data is grouped, Vertica applies over twelve compression schemas. Dependent on the data, the system chooses which to apply. Typically 50-90% compression is achieved. The unique ability to query data directly in encoded form requires fewer CPU operations. This results in storing more data, providing more views, and faster results with less hardware.
- **Clustering**—Because real-time analytics makes log-based recovery impractical, columns are duplicated so that if one node fails, a copy is available. The new hardware queries the rest of the cluster for data it needs and rebuilds missing objects from the other nodes. This also enables the system to scale-out infinitely by simply adding more hardware.
- **Continuous performance**—Concurrent loading and querying enables real-time views and eliminates nightly 'load windows'. On-the-fly schema changes can be made to add columns or projections without downtime. Automatic data replication, failover and recovery—'active' redundancy increases performance as nodes will recover automatically by querying the system. Unique ability to load/insert data at very high rates of speed on a 24x7 basis with zero administration while delivering blinding query performance against large volumes of data continuously.



“Vertica has rapidly gained enterprise customers with its high-performance EDW architecture. With an increasing focus on performance, scalability, optimized storage, and in-database analytics, Vertica Systems positions its EDW offerings as a robust platform for the most demanding enterprise analytics. ... Vertica’s customer momentum, coupled with its focus on enhancing its columnar-based EDW architecture, gives it a competitive advantage.”

— James Kobielus, February 10, 2011
Forrester Research, Inc.

Designed with Converged Infrastructure advantages

The HP Vertica Analytics System provides you with higher levels of operational efficiency and highly-reliable systems with 2x the load speed of the nearest competitor and query performance that is 50x-1000x faster than any product on the market.

The integrated, turnkey appliances are built with the Converged Infrastructure advantages of a common modular architecture. Quarter-rack, half-rack, and full-rack configurations are optimized with processor core density more than 45% greater than competing solutions. The HP BladeSystem server blades and HP StorageWorks infrastructure deliver capacity optimization up to ~1PB database in a single rack. Non-disruptive growth and transparent scale of performance and capacity can be achieved in multi-rack solutions as well.

Incremental Vertica software licensing allows you to grow your database based on the amount of data loaded, not by requiring you to buy additional DBMS licenses every time CPUs are added to the system.

Delivering immediate value

Beginning with HP Factory Express Services, On-site Installation and Startup Services, this highly-reliable system includes one year of HP Support Plus 24, 24x7 hardware and software support for the entire solution with a single point of accountability.

HP offers a variety of additional support services to meet your business needs, including configuration and professional services.

Once operational, efficiency is improved through HP Insight and HP integrated Lights Out 3 (iLO) management software. These solutions are designed to simplify server installation, configuration, and maintenance throughout the entire server lifecycle.

Easy to deploy and maintain because of the automatic physical design with integrated Database Designer, the HP Vertica Analytics System is ready to roll in days, not months, eliminating many time-consuming DBA duties.

Harnessing the power of real-time analytics

The HP Vertica Analytics System delivers radically improved database price-performance allowing more data to be loaded faster to uncover insight into your ‘Big Data’.

- Low-latency “in-time” analytics
- Up to 10x better load performance
- Hybrid in-memory/on-disk architecture
- Linear scalability and full parallelism

More questions, more data, faster answers

Consume, analyze, and make informed decisions at the speed of business with the HP Vertica Analytics System. Delivering the fastest return on information makes it possible for you to monetize your ‘Big Data’ with real-time analytics in a matter of minutes, not days.

For more information

To read more about Vertica go to www.vertica.com.



Get connected

www.hp.com/go/getconnected

Get the insider view on tech trends, alerts, and HP solutions for better business outcomes

