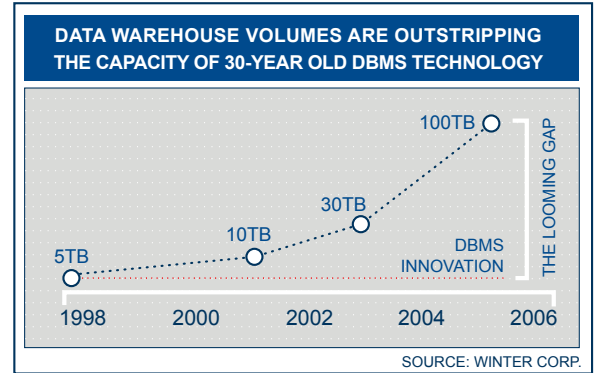


THE TIME FOR DBMS INNOVATION IS NOW...

Regulatory compliance, increased competition and other pressures have created an insatiable need to accumulate and analyze larger and larger quantities of data. Over the last decade, the largest data warehouses have increased from 5 to 100 terabytes (source: Winter Corp), and by 2010, TDWI estimates that most of today's data warehouses will be 10 times larger.

Yet, during the same period, there has been little database management system (DBMS) innovation to keep pace. Performing ad hoc queries on such large data volumes does not come naturally for existing DBMSs, which use a row-oriented design optimized for write-intensive transaction processing workloads rather than for read-intensive analytical workloads. Desperate for better performance, row-oriented DBMS customers spend millions of dollars annually on stop-gap measures such as adding DBA resources, creating and maintaining OLAP cubes or replacing their DBMS with expensive, proprietary data warehouse appliance hardware.



THE VERTICA DATABASE™ ADVANTAGE

At Vertica we believe it's time for significant innovation in the database industry. Led by database research pioneer Michael Stonebraker, we have invented a brand-new RDBMS architecture that provides blinding query performance for databases scaling from hundreds of gigabytes to hundreds of terabytes and where end-user requirements change rapidly. Here is what sets Vertica apart:

\$ Radically Improved Database Price-performance

Benchmarks show Vertica is 20 to 100 times faster than traditional RDBMSs and data warehouse appliances.

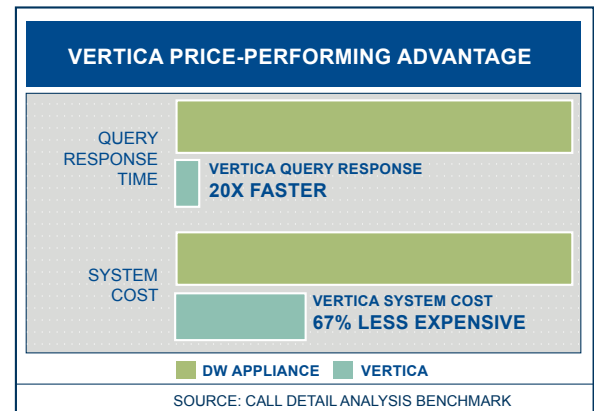
Vertica runs on commodity hardware; couple that with its unique pricing approach and you can deploy large-scale query-intensive databases for a fraction of the costs of other solutions.

□ Painless Scalability

Vertica runs on standard, commodity hardware running Linux and is optimized for cost-efficient distributed computing environments such as grids and clusters. In addition, Vertica licensing is

based on the amount of data you store—not your database hardware configuration. Coupled with the Vertica price/performance advantage, this means you can scale the size and usage of your database without:

- Disruptive database platform changes
- Proprietary data warehouse appliance hardware
- Paying incremental DBMS software licensing fees each time you add CPUs to the system





DBA Liberation

We've built a lot of DBA "know-how" into Vertica to keep it running efficiently without much administrative overhead. High availability, disaster recovery, schema design and physical optimization are performed automatically, freeing your real-world DBAs to focus on higher-value-added activities.

WHAT MAKES VERTICA DATABASE™ UNIQUE – KEY INNOVATIONS

From a database developer perspective, Vertica Database™ looks pretty standard; it supports SQL, ACID transactions, JDBC and works with popular ETL and BI reporting products. Underneath the covers, it's a different story. The Vertica Database is designed to aggressively economize disk I/O and is written natively to support grid computing. Vertica is a 21st-century solution for today's large-scale, read-intensive database applications, featuring ground-breaking architectural features such as:



Column store architecture – 20x - 100x faster query performance

In a row-oriented DBMS, row values are stored contiguously, and to process a query, the DBMS must read data from every row and column (even columns not specified in the query). Bitmap indices, data cubes, materialized views, and other database exotica help, but for many databases, managing data structures to optimize performance for every query and use is prohibitively complex, and these data structures often impose dramatic storage space and update performance overhead. Vertica stores the values for each column contiguously meaning that it only needs to read the columns being queried. This approach dramatically improves query performance by eliminating unnecessary disk and memory I/O.



Heavy use of compression – Store data in 90% less table space

CPUs are getting faster at a much greater rate than disk bandwidth is increasing, so Vertica replaces slower disk I/O with faster CPU cycles to encode data elements into a more compact form and query them. Vertica's innovative query engine operates directly on compressed data, meaning that it can actually require fewer CPU operations to process the compressed version of a table.

VERTICA COMPRESSION METHODS		
COLUMN CONTENTS	FEW VALUES	MANY VALUES
Sequential	RUN-LENGTH ENCODED	DELTA ENCODED
Non-Sequential	BITMAP PER VALUE	CONVENTIONAL GZIP

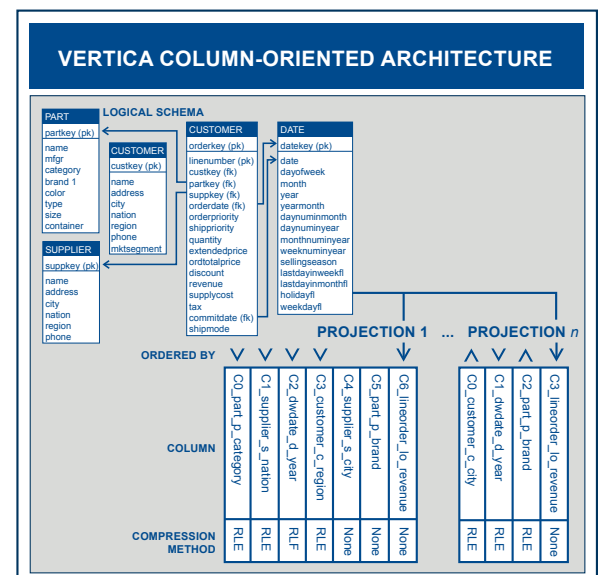


Redundancy & multiple sort orders – Natively designed for Grid computing

Vertica supports logical relational models; physically, it stores data as "projections" – collections of sorted columns (similar to materialized views). Multiple projections stored on networked, shared-nothing machines ("sites") can contain overlapping subsets of columns with different sort orders to ensure high availability and enhance performance by executing queries against the projection(s) with the most appropriate columns and sort orders.

Physically stored as multiple "projections"

- Collections of columns with multiple sort orders
- Distributed across multiple nodes
- Aggressively compressed
- Queries & loading parallelized across them
- Built-in redundancy for high availability
- Designed automatically





Automatic database design – Optimizes for high performance and high availability

Based on DBA-provided logical schema definitions and SQL queries, Vertica automatically determines what projections to construct and where to store them to optimize query database performance and high availability.




Recovery by query – New level of high availability database efficiency

Rather than having a mirrored database backup sitting idle for failover purposes, Vertica leverages the redundancy built into the database’s projections. It queries projections not only to handle user requests, but also for rebuilding the data in a recently restored projection or site. The Database Designer builds the necessary redundancy into the projections it creates such that a DBA-specified number of site failures can occur without compromising the system. Vertica’s approach to recovery avoids bogging down database performance with expensive logging and two-phase commit operations.



Hybrid architecture – Load more data and deliver more real-time answers daily

Data warehouses are often queried by day and bulk-loaded by night. The problem is, there’s too much data to load at night and users are demanding more real-time data. Vertica features a hybrid architecture that allows querying and loading to occur in parallel across multiple projections. Each Vertica site contains a memory-resident Write-Optimized Store (WOS) for recording inserts, updates and deletes and a Read-Optimized Store (ROS) for handling queries. WOS contents are continuously moved into the associated ROS asynchronously. Lightweight transaction management prevents database reads and writes from conflicting so queries can run against data in the ROS, WOS or in both.



VERTICA ENABLES MORE PEOPLE TO QUERY MORE DATA FASTER

Vertica’s radically improved database price/performance, scalability, availability and ease of administration let customers deliver more data to more people for more uses, at less cost and effort. Vertica is ideal for all kinds and sizes of businesses, and in applications such as:

- ✓ Data Marts
 - ✓ Business Intelligence
 - ✓ Data Warehousing
 - ✓ Click Stream Analysis
 - ✓ Fraud Detection
 - ✓ Compliance Reporting
- ✓ Call Detail Analytics
 - ✓ Tick-Store Query Apps
 - ✓ Buying Pattern Analysis
 - ✓ Basel II Compliance
 - ✓ Sales Dashboards
 - ✓ And others

ABOUT VERTICA SYSTEMS

Vertica Systems is the innovator in the market for affordable, high-performance database management systems. Co-founded by database research and technology pioneer Dr. Michael Stonebraker, Vertica has developed remarkably fast database technology for storing and querying very large databases using industry standard methods (SQL). With its unmatched affordability, scalability and flexibility, Vertica’s solution will bring corporate database architectures into the 21st century.