



Revolutionizing Data Warehousing in Telecom with the Vertica[®] Analytic Database

A DBMS architecture that takes CDR, SNMP, IPDR and other telecom data warehouses to the next level of performance, simplicity and cost efficiency.

Copyright Vertica Systems Inc. March 2010

Table of Contents

Table of Contents.....	2
“Where is money leaking from the system?”.....	3
The Time for Database Innovation in Telecom is Now	4
The Vertica Analytic Database Advantage	4
Vertica Analytic Database – Key innovations	5
Column Store Architecture	5
Extensive Compression.....	6
Multiple “Projections” Stored	6
Shared-nothing Parallelism	6
K-Safe based Availability.....	7
Automatic Physical Database Design	7
Hybrid Storage Model.....	8
Deployment flexibility.....	8
In Summary: A DBMS Architecture Built for Telecom Analytics	9
Vertica Customer Experiences	9
NetworkIP.....	9
Comcast	10
Additional Resources	11
About Vertica Systems	11

“Where is money leaking from the system?”

Revenue assurance analysts at a top-tier US-based carrier studied this every day. Primarily focused on detecting fraud, revenue sharing contract violations and incomplete revenue collections, they had the need to query and analyze call detail record (CDR) databases that grow by millions of new CDRs every day. The business angles they examine include:

Customer Analysis	<ul style="list-style-type: none"> ▪ Usage profiles ▪ Aggregate reports for customers with multiple accounts/services ▪ Customer segmentation reports ▪ Customer retention and acquisition
Product Management	<ul style="list-style-type: none"> ▪ Profit margin reports ▪ Analysis of bundle profits ▪ Competitive analysis ▪ Reports of promotion ROI
Revenue Reporting	<ul style="list-style-type: none"> ▪ Lost revenue reporting (under-billing, etc.) ▪ Verifying billing data with call data from switches ▪ Verifying customer orders and billing
Network Cost Analysis	<ul style="list-style-type: none"> ▪ Support data for negotiating access to competitors' networks ▪ Calculating most cost-effective network routes ▪ Network utilization versus capacity reports

The problem was their aging CDR data warehouse:

- Queries took tens of minutes or even hours to answer
- They only had access to a few months of CDR history
- They could not perform ad-hoc analysis
- Annual costs were very high (DBMS maintenance fees, 4 DBAs, data center costs for the SMP server and SAN on which it ran)
- The data warehouse could not be changed easily to support new requirements from legal, regulatory, and engineering or to support wireless or IP data

This caused the company to seek alternatives to their data warehouse DBMS, and it led them to Vertica, which they implemented. The grid-based, column-oriented Vertica Analytic Database is a relational SQL database built specifically to handle today's analytic workloads in telecom. Vertica enables Comcast, Vonage, Level 3, Verizon and others to manage terabytes of telecom data (CDR, SNMP, IPDR, et al) faster, more reliably and more economically than with any other database by providing the following benefits:

- **Speed** – Queries 50x-200x faster by reducing disk IO, and loads at world record speed¹
- **Savings** – Reduces hardware, data center and DBA costs by up to 90%
- **Simplicity** – Automates database design, tuning, replication, failover and recovery
- **Scalability** – Scales out infinitely on “green” grids of inexpensive servers

¹ Vertica teamed with HP and Syncsort in December 2008 to break the data warehouse loading world record—5TB loaded in under 1 hour. Visit www.ETLWorldRecord.com to learn more about the benchmark

This paper describes the innovative architecture of the Vertica Analytic Database and how it's able to provide such remarkable performance compared to other databases.

The Time for Database Innovation in Telecom is Now

The number of subscribers to mobile, fixed-line and cable communications services is growing by millions of people every year, and the volume of CDR, SNMP and IPDR data that communications companies must store and analyze is also exploding, by terabytes per year.

Yet, during the last 30 years, 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 hardware.

The Vertica Analytic Database Advantage

In contrast to OLTP databases like Oracle, SQLServer or MySQL, Vertica has been specifically designed to support analytic workloads, which primarily comprise querying (often ad-hoc) and the insertion of new data (in bulk or on a constant "trickling" basis for real-time analytics). By featuring an innovative architecture that drastically reduces disk I/O, Vertica delivers orders of magnitude faster performance and much lower cost of ownership.

The following performance benchmark results conducted by a Vertica telecom customer (based on their real-life data set) helps illustrate the Vertica advantage:

Vertica reduced query times from hours to minutes and hardware costs by over \$500,000 USD

CDR Data Warehouse	Vertica® Analytic Database	Row-Store DBMS	Vertica Advantage
Benchmark Data Set	1.2 TB Call Detail Records		
Production Data Set	50TB Call Detail Records		
Benchmark Hardware	3 Dell servers – 2 x dual core Opteron 2220, 16GB RAM, 1TB local disk	24-CPU SMP server, large SAN	Over \$500K hardware savings
Mean Query Time (4 queries)	5.4 minutes	2 hours	22x faster
Load Time	5.5 hours	82 hours	15x faster
Database Size (1.2TB Raw Data)	220GB	4TB	18x less storage (82% compression)

The remainder of the paper describes the Vertica architecture and how these performance improvements are achieved. It also explains how Comcast and NetworkIP are using the Vertica Analytic Database.

Vertica Analytic Database – Key innovations

From a database developer perspective, the Vertica Analytic Database looks very standard; it supports SQL, ACID transactions, JDBC, ODBC and works with popular ETL and BI reporting products. Underneath the covers, it's a different story. Vertica is designed to aggressively economize disk I/O and is written natively to support grid computing. Vertica is a 21st-century solution designed to capture and query terabytes of CDR, SNMP or IPDR data and other large, fast-growing volumes of telecom data. Vertica includes ground-breaking features such as:

Column Store Architecture

In Vertica, data for each column is independently stored in contiguous blocks on disk. Column values are linked into rows implicitly, based on the relative position of values in each column. Unlike most database products where all columns for each row are stored together, the Vertica Database only needs to retrieve those columns needed for a specific query, rather than all columns in the selected rows (see *Figure 2*).

Vertica's vertical partitioning approach produces dramatic I/O savings for the large majority of decision support queries that only retrieve a subset of columns. For example, one telecommunications company has call detail records with 230 columns. Most queries touch fewer than 10 of these. In fact, as many as 200 of the columns are rarely used by any query.

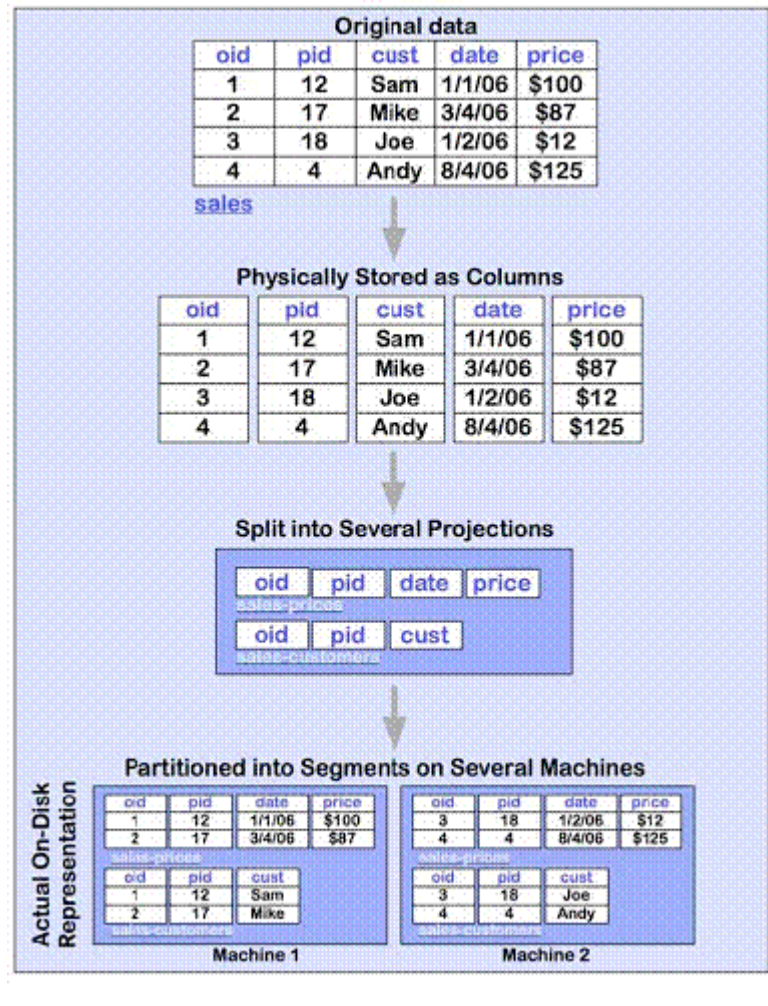


Figure 2: Vertica Analytic Database Architecture

Extensive Compression

The Vertica Analytic Database employs multiple compression algorithms, depending on data type, cardinality, and sort order, to minimize the space occupied by a column. These include run length encoding, delta value encoding, integer packing for integer data, block-based dictionary encoding for character data, and Lempel-Ziv compression, among others. Vertica automatically chooses a good algorithm for compressing data in each column, based on a sample of the data. Compressing data by column often improves compression ratios because the data shares a common data type and value range. Run-length encoding (RLE), for example works best for columns of ordered data, or data with few distinct values compared to the number of rows. This ensures long runs of identical values (such as columns containing dates, gender or state codes), which RLE compresses quite well. Database compression ratios for CDR and SNMP data typically ranges from 8x to 13x relative to the size of the ASCII input data.

The Vertica Analytic Database can also directly access data values in compressed form. This feature avoids the need to spend processor cycles expanding column values during query execution and sets Vertica apart from other database systems that support compression.

Multiple “Projections” Stored

Instead of storing data in base tables as defined in the logical schema, Vertica physically stores multiple views of the table data, called “projections” (similar in concept to materialized views). Each projection contains a subset of the columns of a table in a particular sort order. Rows in a projection consist of the value at the same position in each of the column stores comprising the projection (see *Figure 2*). Projections can also contain columns from multiple tables, thus materializing joins. To support ad hoc queries, every data element is guaranteed to appear in at least one projection. Vertica automatically selects appropriate projections to optimize query performance for the expected workload. Benefiting from large storage savings due to its extensive use of compression, Vertica can maintain multiple, projections, with different, and often overlapping sets of columns, in several different sort orders (and without the added overhead of storing base tables), to improve performance for a wide range of queries, including ad hoc queries needed for exploratory analysis.

Shared-nothing Parallelism

The Vertica Analytic Database is a shared-nothing, massively parallel processing (MPP) database system, designed to run on a grid of Linux servers connected by a TCP/IP network. Nodes contain commodity, multi-core processors with 2 to 4 GB of RAM per core. Storage can be directly attached to each node, or can be SAN-based. Vertica partitions data across the cluster automatically. For example, in the case where a star or snowflake database design, but its not required by Vertica, has been implemented, fact tables are hash segmented across the nodes of the cluster (see *Figure 2*), and dimension tables are typically replicated on each node in the cluster. Very large dimensions are segmented on the same key as the fact table. This limits the need to share data among nodes during query execution. Queries can be initiated on any node. The query planner determines what work needs to be done to answer the query, distributes it to participating nodes, collects each node’s partial result and prepares the final answer to be sent to the requestor.

K-Safe based Availability

The Vertica Analytic Database maintains multiple stored projections, which can also serve as redundant copies of the data for purposes of high availability. By imposing an additional constraint, namely, that the system guarantees that projections are segmented such that each data element exists on multiple nodes, Vertica implements intelligent data mirroring as an integral part of the database. Vertica calls this K-Safety, where k is the number of node failures that a given set of Vertica projections will tolerate. Vertica guarantees K-Safety by building $k+1$ replicas of all segmented projections—where each replica has the same columns and partitioning key, though the sort order may differ—and offsetting the distribution of segments across nodes for each replica. K-Safety allows requests for data owned by failed nodes to be satisfied by existing projections on surviving nodes, even though the optimal projection may no longer be available. For any given query, one node processes two segments of data—its own, and the one owned by the failed node—so, generally speaking, in a busy system Vertica may take longer to complete the query following a node failure. For multiple queries, the query planner can distribute the extra work to different nodes, so surviving nodes share the burden. Once the failed node is restored, Vertica uses the projections on the other nodes to automatically re-populate data on the previously failed node.

Automatic Physical Database Design

With the Vertica Analytic Database, users need only specify a logical database schema. Given a logical schema, the Vertica DB Designer automatically generates an appropriate physical database design based on that schema, a sample of representative data and queries, and a space budget for the database (see *Figure 3*). The DB Designer guarantees that any valid query on the schema can be answered by ensuring that all data appears in at least one projection. The DB Designer chooses a compression technique to use for each column. It determines which projections to build in order to optimize performance of the sample workload, which columns and joins they will contain, and how each projection should be sorted. It selects the appropriate segmentation key for each projection and guarantees that the projections satisfy the specified K-Safety level. By taking over the details of the physical design, Vertica simplifies database implementation and allows database designers and administrators to focus on the best logical data model to meet their business needs.

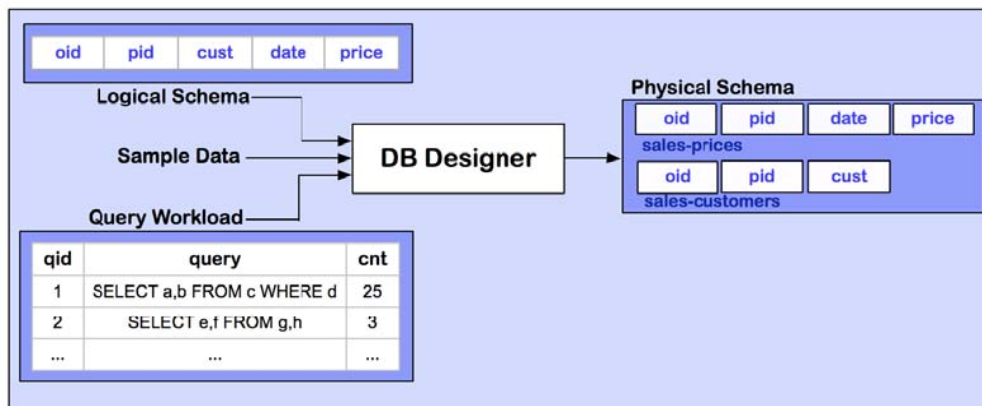


Figure 3: Vertica DB Designer Process

Hybrid Storage Model

A Vertica database caches all database updates to a queryable main memory cache called the Write-optimized Store (WOS) (see *Figure 4*). The WOS organizes data into projections that are stored as collections of uncompressed, unsorted column arrays, maintained in update order. An asynchronous background process on each node, called the Tuple Mover, migrates recent updates to permanent disk storage in the Read-optimized Store (ROS). Data in the ROS is sorted, compressed, and densely packed into variable-length disk blocks that are optimized for query performance.

The Vertica Analytic Database supports snapshot isolation for query processing, which means that queries and updates do not interfere with one another and that read-only queries do not require locking. Updates are collected in time-based buckets called epochs. New updates are grouped in the current epoch until the transaction is committed. Data in older epochs is available for query processing and eventual migration by the Tuple Mover to the ROS (see *Figure 4*).

Vertica's hybrid storage model supports both bulk loads and trickle-feed updates. In Vertica, the Tuple Mover is working in the background to drain the WOS and merge updates into the ROS to keep it current. While clearly designed for read-mostly application, this approach also works for near-real-time data warehouses with high append data volumes. CDRs and SNMPs can usually be seen within seconds of the call or network event.

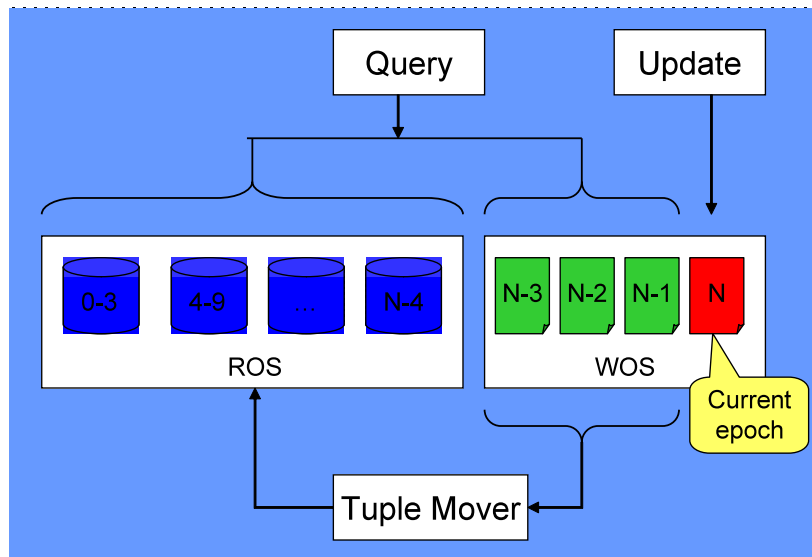


Figure 4: Vertica's Hybrid Storage Model

Deployment flexibility

The Vertica Analytic Database can be downloaded, installed on a cluster of Linux servers and run as described above. It can also be licensed and used on a monthly basis hosted in the Amazon Elastic Compute Cloud (Amazon EC2). The Vertica architecture lends itself well to virtualization, which permits it to be used as a cloud-hosted service. The flexibility of being able to use Vertica for the Cloud enables organizations to have data warehouses provisioned rapidly and without having to budget and procure DBMS hardware and software. This opens up new data warehousing opportunities to support proof of concept projects, short-term analyses (e.g., a 3-month re-pricing project) and seasonal or start-up SaaS ventures that make large volumes of data available to customers for analysis on line.

In Summary: A DBMS Architecture Built for Telecom Analytics

Together, the key features of the Vertica Analytic Database create an elegant architecture for high-speed, cost-effective analysis of large-volumes of CDR, SNMP and IPDR data. Vertica's vertical partitioning by column, extensive use of compression, and hybrid storage model reduce the IO required to execute queries. Telecom records such as CDRs contain many columns per customer, but individual reports cull only a few of these. The types of analytic applications that use a small, arbitrary subset of columns per query are both common, and ideally suited for the vertical partitioning provided by the Vertica Analytic Database. At the same time, Vertica's data partitioning, which divides work across multiple nodes in a computer cluster, supports the very large data volumes to which these applications often grow. These analytic applications also typically support users from many disciplines, interrogating the database from multiple individual perspectives. Here, Vertica's ability to keep multiple physical projections of the data is a natural fit for such usage patterns. In summary, the Vertica architecture was designed specifically to handle the common characteristics of analytics applications in telecommunications.

Vertica Customer Experiences

Vertica, first released commercially in 2007, now has a rapidly growing base of customers in telecom as well as in financial services, analytic applications software, online business services, health care and other industries. Read on to learn how NetworkIP and Comcast are using Vertica:

NetworkIP

NetworkIP is one representative customer – a provider of prepaid calling services to 120 telecom carriers around the world. In its Longview, Texas, data center NetworkIP services 400 million end users and manages billions of call detail records (CDRs). In addition to calling services, NetworkIP offers its customers business management tools in the form of online application software that can be used to analyze CDRs for rate plan profitability, traffic patterns and quality of service. Customers can also change their rate plans online using software provided by NetworkIP.

By early 2007, NetworkIP was outgrowing its data warehouse platform. Service level agreements (SLAs) required that queries against the CDRs be satisfied in 10 seconds, but the database was now so large – at 1.4 TB of data -- that some queries required 45 to 60 seconds. In addition, thousands of summary tables were maintained in the data warehouse and queries were limited to the summary tables, because query against the detailed data with the existing platform would have required inordinate time and resources. For similar reasons, only the most recent 90 days of CDRs were available for online analysis.

Finally, customers wanted to be able to change a calling plan and analyzing the resulting changes in calling behavior the same day. With NetworkIP's existing data warehouse platform, new CDRs were not available until the following day.

NetworkIP conducted an evaluation of data warehouse platforms, considering architecture, performance, scalability and cost, culminating in a benchmark in the last quarter of 2007. In the benchmark and evaluation, Vertica satisfied all of NetworkIP's objectives, outperforming other platforms by a large margin. Vertica was able to satisfy all queries in under 10 seconds, regardless of whether the query examined 30 days, 90 days or six months of CDRs.

With Vertica, queries against the detail data could be supported and many fewer summary tables needed to be built. New CDR data was available for analysis within 30 minutes after the calls occurred, rather than the following day. NetworkIP also uses its data warehouse to help manage

the operations of its own sizable network. With the prior platform, a calculation of the network cost took 12 hours and was therefore done only a once a month. With Vertica, the same cost calculation takes 10 seconds and is performed daily, enabling 365 “course corrections” a year with respect to operating cost.

Vertica was able to support the entire database on a four node cluster of Dell 2950s, providing both increased availability and greatly increased performance, in comparison to the existing operation. Vertica was also far lower in total cost than other options considered.

Comcast

Comcast is the largest cable communications company in the US with over 24 million cable customers and millions of customers for its high-speed internet and telephone services. Among the activities Comcast manages for its customer are over 3 billion ON DEMAND views of movies and television programs each year. The Comcast cable network has millions of electronic and software components that must interoperate continuously, around the clock, every day of the year. Billions of metrics are generated that could indicate a service interruption or other problem. To successfully manage its network, Comcast must proactively collect and analyze this data.

Most of the data involved in monitoring and managing the Comcast network can be characterized as time series data. A typical 4- or 8-byte sample consists of a node identifier, a metric, a timestamp and a value. For example, a sample representing the status of the set top boxes in a particular household might indicate at 8:30pm on November 6 that two TV sets in that household were turned on. Such a status message would be sent periodically by the network equipment – say, once a minute – and the set of all such messages in a day would comprise a time series. In network management, such data is written once to a database and looked at many times in the course of analysis.

Comcast created its own solution to the problem of managing, accessing and analyzing this data, employing open source tools. While this approach worked well for a quite a while, the company recently concluded that the enormous growth in its services and its network management data was resulting in cost, performance, maintainability and access problems.

In 2007, Comcast undertook a search for a data management solution that would meet its demanding requirements, including: (a) a load speed of at least 50,000 samples/second; (b) query response time of two seconds or less; (c) ability to scale to at least 40 TB of data on standard hardware; (d) 10-to-1 data compression; (e) high availability; (f) cost effectiveness; and, (g) ability to provide both detailed data and periodic summaries rapidly in response to queries.

In an evaluation, Comcast found that Vertica exceeded these requirements. Vertica provided sub-second query response – more than 10 times faster than the closest alternative. Vertica exceeded the required load rate by more than a factor of two, while also exceeding the required 10-to-1 compression ratio. In addition, Vertica demonstrated that it could pass a substantial series of fault/recovery tests during evaluation and that it could readily maintain replicas across multiple sites over a WAN. Finally, by running on a grid of inexpensive servers, Vertica delivered the lowest TCO of any alternative evaluated.

When Comcast implemented Vertica, they found that implementation was rapid and readily accomplished. Vertica worked with their existing BI tools, including reporting; desktop dashboard; and, in house ETL. Once implemented, the Vertica-based system gave business users faster access to much more data at much lower cost than the prior solution.

Additional Resources

The Vertica Analytic Database supports SQL and integrates with 3rd-party ETL, analytic and BI reporting tools and applications via JDBC, ODBC and specific language bindings. Therefore, using all your existing SQL knowledge and technology, a Vertica database can be very quickly created and loaded with data.

If you would like to learn more about the Vertica Database or if you would like to evaluate it yourself, then visit the following links:

Telecom Solution Site	www.vertica.com/telecom	Vertica technical and customer resources for telecommunications companies
Resource Library	www.vertica.com/resourcelibrary	White papers, demos, webcasts, system requirements
Vertica Benchmarks	www.vertica.com/benchmarks	See customer-submitted cost and performance comparisons between Vertica and other databases
	www.etlworldrecord.com	Vertica was used to set the data warehouse data loading world record – 5.4TB in 57 minutes
Vertica Customers	www.vertica.com/customers	See who's using Vertica
Evaluate Vertica	www.vertica.com/download	Request a free evaluation copy of the Vertica Analytic Database to download and install

About Vertica Systems

Vertica Systems is the market innovator for high-performance analytic database management systems that run on industry-standard hardware. Co-founded by database pioneer Dr. Michael Stonebraker, Vertica has developed grid-based, column-oriented analytic database technology that lets companies of any size store and query very large databases orders of magnitude faster and more affordably than other solutions. The Vertica Analytic Database's unmatched speed, scalability, flexibility and ease of use helps customers like JP Morgan Chase, Verizon, Mozilla, Comcast, Level 3 Communications and Vonage capitalize on business opportunities in real time. For more information, visit the company's Web site at <http://www.vertica.com>.