

Dialing Up Blazingly Fast Call Detail Analysis with Vertica

Founded in 1998, NetworkIP is the innovator and leader in providing prepaid services to the telecommunications industry. NetworkIP provides a carrier-grade platform, products and services that enable its customers to offer telephone calling cards, audio conferencing, and other prepaid services on a turnkey basis. The company's customers include major carriers, telecom service providers, distributors, resellers, CLECs, ILECs, ISPs, and publishers.

All these customers want the same thing: the highest quality of service for their calls, at the lowest possible price. NetworkIP achieves this with the help of its award-winning Integrated Connection Solution (ICS) platform; this technology enables rapid, low-cost deployment of converged voice, data and IP solutions over wired, wireless and IP networks.

The NetworkIP ICS Platform

To further differentiate itself from other providers, NetworkIP puts a lot of emphasis on customer service and a partnership approach. Customers rely heavily on access to the data in NetworkIP's call detail records (CDRs) in order to make business decisions such as calling program rate changes. NetworkIP customers are experts in marketing and selling via their distribution channels, but they lack the IT infrastructure and expertise for business intelligence.

NetworkIP answers this need by using advanced technology to provide the analytical tools that its customers need to run their businesses. The company built the ICS platform, which comprises a data warehouse and a web-based data-mining application. The application connects with multiple databases and provides reports to customers as well as to internal employees. Using the ICS application, customers can mine the CDRs to analyze their traffic patterns, monitor and forecast their profitability, and perform other analyses.

In the six years since NetworkIP deployed the data warehouse, Director of Business Application Development Stephan Broquie and his team have built thousands of summary tables to support all the different reports that customers want: for example, to slice data by month, by year, by quarter and so on. The ICS platform creates the original CDRs. It then sends them to a summarizing engine, which is based on MySQL; the summarizing engine breaks the CDRs up into the various tables.

The Vertica Solution At a Glance

The Customer



www.networkip.net

The Industry

Telecommunications

The Application

Data warehouse and web-based BI application for analyzing CDRs

The Benefits

- Queries up to 1000X faster
- Eliminates the need for aggregate data and enables faster, more granular analysis
- Cost-effectively store and query two years' worth of CDRs (vs. 90 days')
- Near real-time analysis (30 minute vs. one-day lag time)
- Network cost analysis query time reduced from 12 hours to 10 seconds—enables daily cost adjustments vs. monthly

The Need for a New Analytic Database within the ICS Platform

The application worked very well in the beginning, and customers loved it. But as NetworkIP's business has evolved, its customers have become savvier and more data-hungry. They started requesting reports that Broquie and his team couldn't support easily with summary tables. In short: the ICS had become a victim of its own success.

Broquie explains: "Our customers' focus used to be mostly on what happened last week and last month. Today, customers want to know things over a long period of time. They want to look at trends and at things that happened over the last 12 months, and to be able to compare this year with last year. Most of our queries have only been able to accommodate a date range of 90 days. And, within a 90-day date range, the system will provide a response within a minute or so. This clearly wasn't good enough."

Faced with customer demands that were outpacing the application, Broquie and his team began looking for a database that would let customers directly query the raw CDRs for any period of time. The database would have to let the data warehouse store at least two years' worth of CDRs – or about 1.2 billion records, each containing about 200 pieces of information – and deliver responses in 10 seconds or less for 90% of queries. Broquie and his team, which includes three seasoned database administrators (DBAs), knew from experience that a traditional relational database system (RDBMS) couldn't provide the required performance without making a prohibitively large investment in DBMS server hardware.

The Solution: High-Performance CDR Analysis using the Vertica Analytic Database

In September 2007, NetworkIP's search ended at Vertica Systems. Vertica has a high-performance, column-oriented RDBMS that runs on industry-standard Linux-based hardware, and NetworkIP proved that the Vertica Analytic Database is ideally suited for the application's performance demands and storage requirements. The Vertica Analytic Database organizes data on disk as columns of values from the same attribute, as opposed to storing it as rows of tabular records as in traditional RDBMSs. When a query needs to access only a few of those attributes – typical for NetworkIP customers – it only needs to read those columns. By comparison, with traditional RDBMSs, queries have to read all values in a table, wasting I/O bandwidth and making the queries very slow.

The Vertica Analytic Database also uses compression very aggressively, both of data on disk and on data "in motion" during queries, which further enhances query speed while enabling cost-effective storage management. Companies can store 10 terabytes of data in as little as 1TB of disk space.

Because of its speed, the Vertica Database lets NetworkIP avoid having to work with summary or aggregate data; customers and in-house staff can directly query the raw CDRs. This capability improves accuracy, reduces DBA time, and allows NetworkIP and its customers to do much more granular analysis.

"Vertica gives us near-real-time access to data," says Broquie. "ICS Reports include data that may be only 30 minutes old – instead of a day old. Formerly, I couldn't provide our customers any information for today until tomorrow, because our summarization process took place at night in batch. Vertica's unique trickle loading capability enables data to flow into the system on a 24x7 basis so customers can quickly evaluate the affects of calling plan changes and re-adjust them accordingly—cycles which can now happen several times per day."

7 Key Vertica Database Innovations

1. Column store architecture
2. Aggressive compression
3. Concurrent load and query
4. Automatic database design
5. High availability without hardware redundancy
6. Runs on commodity hardware
7. Scale by adding inexpensive servers

“The bottom line is that Vertica is solving problems. It is solving our current problems and it will allow us to launch a new generation of applications that will help our business grow.”

Stephan Broquie
Application Development Director
NetworkIP

The company is also seeing big performance advantages with internal use. For example, it has cut the time to perform a critical monthly cost analysis of toll-free traffic from 12 hours (using the MySQL engine and summary tables) to 10 seconds (using the Vertica Database). That improvement is game changing—NetworkIP can make 365 cost basis corrections per year to boost profitability now instead of 12. They also analyze call data in near real-time, making routing decisions based on complex trade-offs, to deliver the best customer experience at the lowest cost.

“We spent most of last year creating a new data mining center, with a reporting engine that is very flexible and allows us to create new reports very quickly,” Broquie explains. “What I didn’t have, until now, was a database engine that could support this type of flexible reporting.”

The Vertica Analytic Database also will save NetworkIP money in hardware and administration over the long term. NetworkIP is running the Vertica Analytic Database on a four-node cluster of Dell 2950 servers. Each server has two quad-core CPUs with 16

gigabytes of RAM and four hard drives, and a RAID zero configuration. Each server has 1.2 terabytes of available space. Vertica’s “K-safe” architecture provides built-in redundancy and automatic recovery for high availability.

To scale database performance and storage, Broquie and his team can just add inexpensive, off-the-shelf servers to the Vertica cluster, instead of having to perform a traditional “rip and replace” hardware upgrade.

“Every prepaid service company struggles with back office systems,” Broquie says. “We don’t have millions of dollars to spend on an IT infrastructure but we need to be able to deliver the appropriate information to the customer quickly. Vertica will tremendously help with that.”

Of the working relationship with Vertica, Broquie says: “Vertica is very nimble and flexible and has been very responsive to our needs so far. The company is much like our company, in size and in culture. We are a very technology-oriented company. We move fast and we try to accommodate our customers the best we can, so I saw a lot of synergy between our two companies.”

ICS Futures – Enabling New Analytic Possibilities with Vertica

With Vertica, NetworkIP will be able to create new business intelligence tools because it now has a database engine that can create new CDR-mining reports in a day or so instead of a week or weeks.

Broquie is looking at new reports that will allow the company to monitor its profitability more accurately and adjust its costs faster. He is also planning new reports that will let customers monitor their products better. One example: a loyalty program tracker that would identify frequent buyers of calling cards by tying incoming phone numbers to calling-card numbers and then tracking purchase volume and usage over time. Reports like this would let NetworkIP's customers judge how products are working and adjust their marketing and sales strategies in near-real-time; for example, they could change the rate and see what effect it had on product sales.

"The bottom line is that Vertica is solving problems. It is solving our current problems and it will allow us to launch a new generation of applications that will help our business grow," concludes Broquie.

Try the Vertica Analytic Database Yourself

Getting started with the Vertica Analytic Database is easy. It supports SQL and integrates with ETL, and analytical and reporting tools as well as business intelligence applications via JDBC, ODBC and specific language bindings.

If you would like to learn more about how the Vertica Analytic Database can help your company more effectively perform data analysis please visit www.vertica.com or call +1-978-533-3500 to find out more.