

## Marketing Services Site Leadsplease.Com Gains a Competitive Edge with Vertica

Leadsplease.com is one of a new breed of companies using the Internet to deliver innovative business services to direct and channel marketers. Based in Scottsdale, Arizona, Leadsplease helps companies of all sizes create and purchase customized telemarketing and mailing lists on-line within minutes.

The Leadsplease.com service saves enormous amounts of time and money for companies and brands that sell mortgages, cars, insurance, or other goods and services directly or through franchise, dealer or agency networks.

### The Application

Leadsplease.com is a list-prospecting tool that lets users interactively build mailing lists by defining their target markets by geography and audience demographics. Leadsplease.com incorporates 300 million consumers, businesses, homeowners and recent movers provided by industry-leading market data source, Experian. The Web-browser-based application generates SQL queries of the data in response to the user's geographic, demographic and firmographic selection criteria and returns list counts and the cost to purchase and download the list. When the list comprises the type of contacts and is a size that fits the user's budget, they purchase the list on-line through Leadsplease.com and download all the contact information to their own computer.

To help ensure fast query speed for the list counts, the data resides in a data mart on Leadsplease.com's servers and is updated and cleansed weekly, under special arrangement with Experian. The database is nearly half a terabyte (~500GB) and has a very wide data table structure with more than 300 columns of demographic information and 300 million rows.

### The Problem

The Leadsplease.com list-prospecting application has been very successful and popular with customers, and Leadsplease.com continually strives to improve the customer experience by improving the performance of the application. However, improving performance has been difficult because of the limitations of the company's relational database management system (RDBMS). There were three problems.

1. Query speeds were slow because the RDBMS had trouble handling the wide data table. As a result, the application was too slow at returning list counts: Customers were getting list counts back in 15-20 seconds. Although 20 seconds may sound fast, it

### The Vertica Solution At a Glance

#### The Customer



[www.leadsplease.com](http://www.leadsplease.com)

#### The Industry

Direct Marketing Services

#### The Application

- Enables companies of all sizes to create telemarketing and mailing lists on line in minutes
- On-line database contains over 300 columns of demographic and firmographic data for 300 million consumers and businesses
- Using Vertica cut query times from 20 seconds to 50 milliseconds
- Runs on a cluster of 2 HP ProLiant DL380 servers
- Self-tuning eliminated indexing and optimization costs
- Sub-second query performance enabled Leadsplease to make the service open to the public

wasn't fast enough for the step-by-step and iterative list-creation process. A direct marketer, for example, would type in his geographic information and have to wait 20 seconds to get a count, then choose his demographic selects and wait another 20 seconds to get another count, and then repeat this process again and again until the target mailing list contained the desired number and type of addressees. The slow query response also meant that the Leadsplease.com application was tied to a traditional click-and-refresh user interface model. This presented an obstacle to Leadsplease.com' vision of delivering a richer, Web 2.0 interface with integrated mapping and interactive graphical charting of results.

2. The data required a lot of manual indexing, which was not only time-consuming and expensive but also resulted in a larger and larger database and longer and longer load times each time the files were updated.
3. As the database grew larger through indexing, Leadsplease.com realized that it would have to make very large investments in hardware, including network storage.

*"The people at Vertica really wedded themselves to our success. We've become really one team, and it's obviously made an enormous difference to the success of the project. In concept, in execution, it's been a delight working with Vertica."*

Graham Brown  
 VP Sales and Marketing  
 Leadsplease.com

In 2008, the company decided that it needed to solve these problems once and for all as it prepared to deploy the next generation of Leadsplease.com, which was designed to give users a more real-time experience.

"We had to look at ways that we could, economically and quickly, manage the data and get a fairly fast result," says Graham Brown, vice president of sales and marketing for Leadsplease.com. "We knew that the speed of the counts would really dictate the level of sophistication that we could offer in the second generation of the application."

Leadsplease.com thus began an extensive search for databases that could meet its criteria for both query speed and database-rebuild speed. The company tested a number of traditional relational databases (including the incumbent, MySQL), but found that none could meet its requirements.

## The Solution

Leadsplease.com began working with Vertica Systems. Vertica has a high-performance, grid-based, column-oriented analytical database that runs on clusters of inexpensive, off-the-shelf Linux-based hardware.

The Vertica® Analytic Database is ideally suited for the application's wide tables. Vertica 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 columns – typical for Leadsplease.com users – Vertica only reads those columns. By comparison, traditional "row-oriented" databases scan all columns in the table regardless of how many columns the query references, which wastes I/O bandwidth and makes the queries very slow. The Vertica Database also uses compression very aggressively, both of data on disk and of data "in motion" during queries, which further enhances query speed.

Using Vertica, Leadsplease.com was able to cut query times against 300 million rows of data from 15 to 20 seconds to less than 1 second – sometimes as little as 50 milliseconds. Vertica is fast at “OR” queries, which are heavily used by customers but had been particularly difficult for the previous RDBMS.

“Because of the speed of Vertica, we’ve been able to combine the steps into one interactive customer experience,” says Brown. “The user can now define the geography and the map will draw. He can do demographic selects – for example, ‘find only people with an income of up to \$100,000 a year’ – and then, in real time, Leadsplease.com will tell him instantly how many consumers or businesses are in that geography with incomes of up to \$100,000 a year.”

The Vertica Database has a shared-nothing, grid-based parallel architecture, which allows the database to scale effectively on clusters of commodity servers – saving money on hardware. Additionally, Vertica’s architecture enables data to load without impeding query operations, which means Leadsplease.com can run 24x7 without downtime during “load windows.”

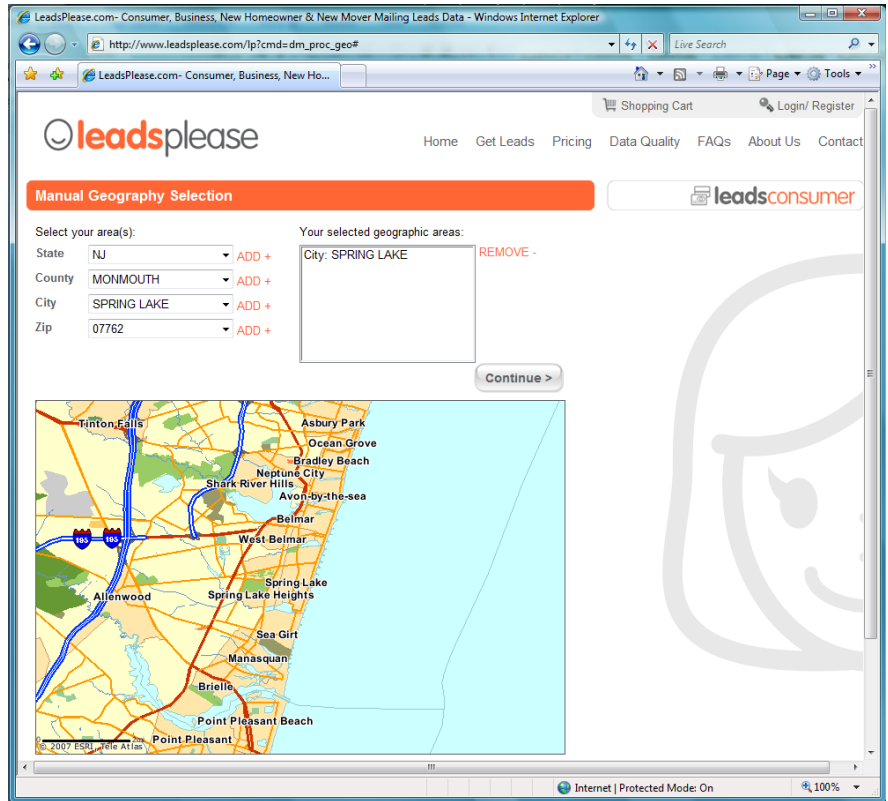
Leadsplease.com can now load data updates more quickly because no indexing is required, thereby preventing “database bloat.” The company is running the database on two very fast multiprocessor servers (eliminating the need for networked-attached storage devices), and expects to be able to manage the application without hiring additional DBAs.

Vertica proved very easy to deploy, enabling Leadsplease.com to build out the database without DBA support. Because the database is accessible using standard SQL and requires no tuning, any developer can load data and start performing queries. Leadsplease.com expects that Vertica’s “load and go” capability will significantly reduce ongoing maintenance costs. By comparison, in its test of traditional DBMSs, IT staff had to spend hours configuring indexes, building materialized views and tuning the database.

## The Future

The advantages provided by the Vertica Analytic Database are essential to the future of Leadsplease.com. Leadsplease.com intends to offer new business intelligence services and Web 2.0 interface enhancements that serve the current customer base as well as enable expansion into new markets.

Customers will be able to upload their own data and analyze it. For example, a cosmetic surgeon could upload his best customers in a file. Within a couple of seconds, Leadsplease.com will be able to create a statistical model of his customers, find his market penetration against a universe of people that look like his existing customers, and present him with maps, pie charts and market penetration analyses that will enable him to zero-in on the best prospects.



## The Final Word

“The Vertica Database enables us, through very fast counts, to radically improve our user interface,” summarizes Brown. “It enables us to reduce the cost of database management. It reduces the amount of storage space that we require on the servers, therefore reducing the cost of hardware. And it radically improves the speed of the refresh of the database. The product is quite radical.”

Of the working relationship with Vertica, Brown says: “It really has been a great experience. They’ve really wedded themselves to our success. We’ve become really one team. And it’s obviously made an enormous difference to the success of the project. In concept, in execution, it’s been a delight working with Vertica.”

## Try the Vertica Analytic Database Yourself

If you would like to learn more about how the Vertica Analytic Database can help your company transform business analytics or to request an evaluation copy, please visit [www.vertica.com](http://www.vertica.com) to find out more.

### 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