

## hMetrix – Revolutionizing Healthcare Analytics with Vertica & Tableau

The cost of healthcare in the US continues to escalate. Consumers, employers, and the government paid out \$2.2 trillion in 2007, or 16.2% of the GDP – \$7,421 per person – for healthcare, according to the U.S. Centers for Medicare and Medicaid Services.

When organizations need hard facts and evidence-based analysis for combating high healthcare costs, many of them turn to George Chalissery and his colleagues at hMetrix. Chalissery is a former software company executive and a veteran of healthcare consulting. He founded hMetrix in 2005 to provide a simpler, more accurate and more efficient way to analyze healthcare data. His approach combines world-class data analysis with extensive healthcare economics expertise and highly personalized service.

hMetrix's clients include employers, the government, healthcare providers, health plans, pharmaceutical companies, disease management firms, benefits consultants, and policy-makers. The company's services help these clients make better decisions faster.

### The Application

The company excels at combining large sets of health care data – such as clinical data from healthcare providers and claims data from the Medicare/Medicaid system and private insurers – and performing complex analyses on the data. Examples include predictive modeling for disease management companies, to identify the highest-risk patients; and comparative analysis of hospitals and physicians for insurers, to enable them to create effective provider contracts. hMetrix also does outsourced analytics and reporting for employers, to help them understand healthcare utilization patterns and properly assess their risk when self-insuring.

hMetrix traditionally produces static reports created by its analysts. However, clients began asking for direct access to the data so they could ask more questions themselves – without having to go through an analyst. Clients appreciated hMetrix' expertise in acquiring, cleaning and analyzing data, but they wanted the ability to do more ad hoc querying of cleaned data.

### The Vertica Solution At a Glance

#### The Customer



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

#### The Industry

Healthcare

#### The Application

- Healthcare business intelligence
- Data warehouse enables complex *ad hoc* querying of large datasets by healthcare professionals
- Point & click, visual data analyses with Tableau enable faster, better and more strategic decisions
- Improved query performance by an average of 70X
- Dramatically reduced hardware costs via Vertica compressing the data by 8X, which also eliminates need for additional disks in order to get I/O performance; runs on commodity hardware
- Rapid integration and easy data migration; integrates with regression analysis, ETL and reporting tools
- Installed and supported without DBAs and IT administrators

## The Problem

hMetrix suddenly needed a way to move its reporting out of the back room and into a more interactive, non-technical user environment. This move presented two challenges.

The first challenge was ensuring extreme ease of use. The solution had to be intuitive for healthcare professionals—as easy as using Microsoft Office—and not require any programming or SQL knowledge.

The second challenge was accommodating the size of the datasets. For example, Medicare/Medicaid datasets can include hundreds of millions of rows, so any solution would have to be able to cost-effectively warehouse and query very large datasets. On the other hand, claims information datasets from private insurers can vary widely in size depending on the study. Therefore, the solution would also have to scale, providing a consistent user experience regardless of dataset size or complexity. The solution would also have to be fast and flexible running over the Web.

As a small company, hMetrix also needed a solution that could be deployed quickly and be operated without a lot of IT overhead. “Most of us are analysts who like to look at data and make sense of it,” explains Chalissery. “We’re not programmers or builders of IT infrastructure. We’re not database experts who can build indexes or optimize the database to run faster. We didn’t want to hire an army of contractors to help us implement. We needed to keep the team as small as possible and things as efficient as possible.”

After evaluating several options, hMetrix chose a combined solution from Vertica and Tableau Software.

## The Solution

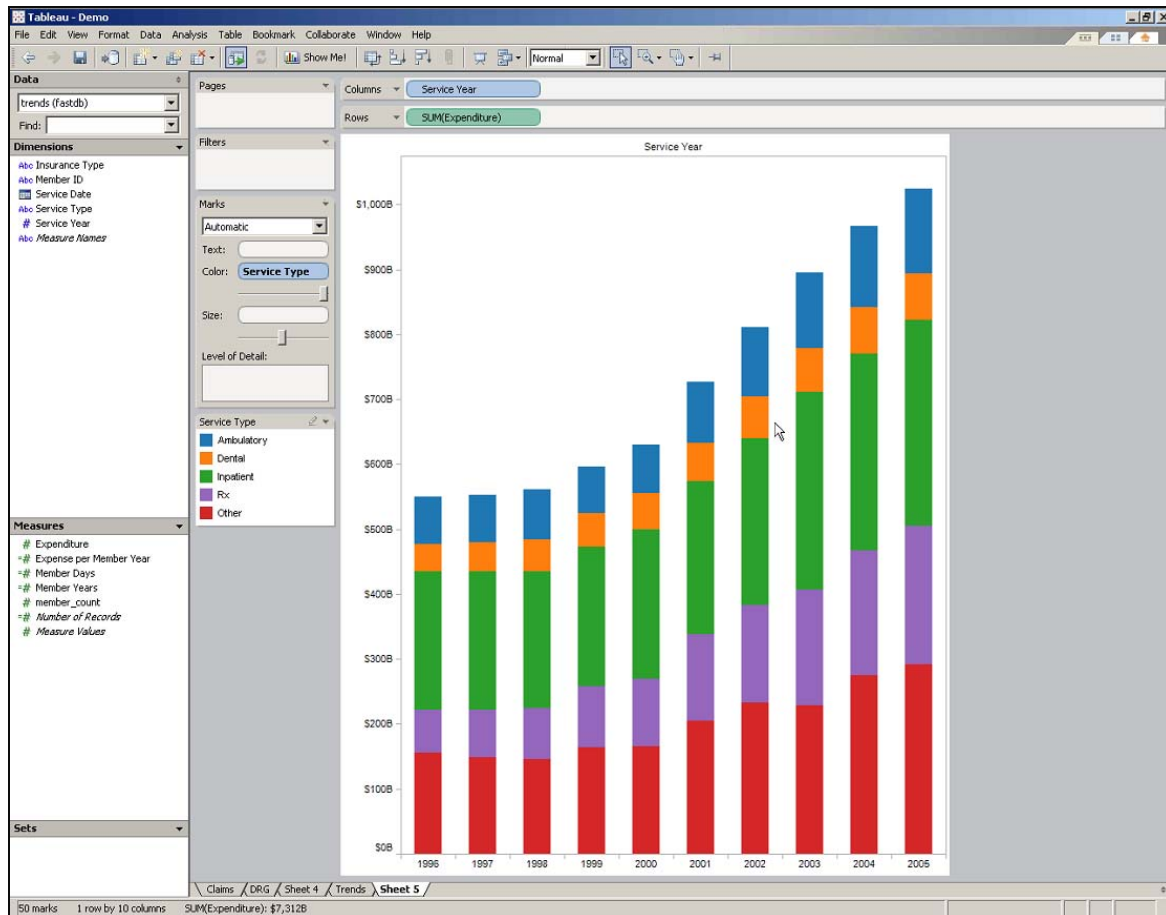
The Vertica Analytic Database is a high-speed, relational SQL database management system (RDBMS) purpose-built for analytics and business intelligence. Tableau provides fast visual data analysis and business intelligence tools that can query back-end databases on the fly using SQL commands.

The Vertica Analytic Database has a column-oriented, MPP architecture, which stores data in the same format by which it is queried – in columns instead of by rows – which makes it very, very fast. By comparison, with traditional, row-store databases, queries must read all values in all rows in a table, wasting I/O bandwidth and making queries very slow. The Vertica Database also spreads the workload across a cluster of inexpensive servers rather than requiring expensive “big iron” to run the data warehouse. Vertica uses compression very aggressively, both of data on disk and of data “in motion” during queries, which further enhances query speed while enabling cost-effective storage management.

*“Tableau and Vertica work very well with datasets that have a large number of dimensions and measures. That’s thanks to the strength of both products. Our day-to-day use of these two products together has been extremely successful.”*

*George Chalissery,  
CEO, hMetrix*

Tableau's software was developed originally at Stanford University to help U.S. Department of Defense's non-technical knowledge workers to rapidly query and understand tremendous amounts of data. The company's solutions – Tableau Desktop and Tableau Server – are as easy to use as Excel but as powerful as query and analysis tools costing thousands of dollars more. People can do fast analytics, visualize data and build dashboards – all without any knowledge of programming. They just point their Tableau desktop at a dataset in the Vertica Analytic Database and go.



hMetrix clients can now quickly test different analyses and instantaneously view the results, often discovering meaningful and useful information that in the past might have remained hidden. The Tableau software leverages the Vertica Analytic Database to do the query processing and calculations, while the desktop and web based client does the interactive manipulation and visual display of results – making analysis very fast and easy. Analyses recalculate as fast as spreadsheets. “Tableau keeps people’s interactions with the technical aspects to a minimum and it gets the answer right,” says Chalissery. “Because Tableau provides easy access to the data, our clients are able to participate in the analytics, often providing insights and knowledge that increase the value of our solution. Tableau was one of the best things that ever happened to us.”

The solution based on both products met hMetrix' requirements for extreme speed, flexibility, ease of use and ease of deployment and administration. The company was able to download both products from the Web, install them, answer a few questions posed by the software, and get them running – all without IT assistance. Tableau Desktop and Tableau Server connect directly with the Vertica Analytic Database without requiring any additional deployment efforts other than licensing and installing Tableau.

Therefore, hMetrix was able to simply transfer data from MySQL into Vertica and connect Tableau to it. The Vertica Database has fast data loading (a 500-million row database loaded in less than 30 minutes) as well as built-in redundancy and automated database administration. Vertica also worked out-of-the-box with hMetrix' other data mining, ETL and reporting tools.

Before purchasing Vertica, hMetrix benchmarked it against the company's current database, MySQL. In the tests, Vertica improved query performance by many orders of magnitude in some cases, with an average improvement of 70X – running on the company's existing hardware configuration.

Query	Vertica	MySQL	Improvement
1	2.5 sec.	283 sec.	113x
2	0.005	202	40,400x
3	6.9	301	44x
4	10.8	1270	118x
...			
154	5.5	357	65x

Figure 1: Results of hMetrix' benchmark comparing query times for Vertica and MySQL

“The results were so astounding that I didn't believe it at first,” recalls Chalissery. “One query on our Medicare database ran so fast – five milliseconds – that I told my team member that he must have misunderstood the units. That he must have meant five seconds. This query used to take a day when we first started running it over a decade ago. In fact, we often ran it overnight, and still took several minutes when we upgraded to more powerful hardware running MySQL. That's how shockingly fast Vertica is.”

hMetrix' standard system is a dual-socket, dual-core x86-64 CPU with 32 GB of RAM and one TB of disk space, running Red Hat Enterprise Linux 5. hMetrix has since found disk space less important because Vertica delivered 8X compression of raw data, reducing the amount of space required to store the data.

Most importantly, the Vertica-Tableau solution enables hMetrix to deliver more value to its clients – value far beyond what used to be possible with its traditional static reports. Customers can further analyze data at any level of detail now on a more ad-hoc basis using Tableau. The data is much more valuable and actionable in helping clients meet their strategic objectives.

For example: for a disease-management firm, hMetrix runs state-of-the-art, customized predictive-modeling algorithms to identify those patients that could most benefit from a targeted management program. hMetrix then runs the resulting predictions back through the ETL process and loads both the detail patient data and predictions into the data warehouse. The next time the client accesses the data, he or she can independently perform the queries and work with information previously understood only by analysts, such as yes/no indicators and probability scores. This was difficult in the past and hMetrix' clients treated the models as black boxes and could do little to peer inside the model to understand its workings using empirical evidence. But now, hMetrix' clients are able to participate in the model building, often providing insight that enhances the accuracy and value of the solution.

### **The Final Word**

Chalissery concludes: "Providing strategic advice to healthcare practitioners and policy makers is very challenging, and there's nothing more persuasive than hard data. Using Tableau and Vertica together gives us and our clients a competitive edge by being able to query data in an ad-hoc manner 70 times faster than with our old database. Tableau visualizes findings in a way that helps clients immediately see paths to better patient care and profitability."

### **Try the Vertica Analytic Database Yourself**

Visit [www.vertica.com](http://www.vertica.com) and [www.tableau.com](http://www.tableau.com) to learn more about how the Vertica Analytic Database and Tableau Software solutions can give your company shockingly fast and impressively easy analysis of huge datasets. You can also download evaluation copies of both products at the web sites.