

Netcore Solutions

Netcore's AI-Engine Raman, built on Vertica, empowers the marketer with real-time insights and helps maximize marketing ROI - powered by blazing-fast analytics and highly accurate AI models.



Raman

Who are Netcore Solutions?

Netcore is a global Martech product company that helps B2C brands create amazing digital customer experiences with a range of products that help in acquisition, engagement, and retention. The first and leading marketing automation and customer engagement platform, Netcore was established in 1997. With decades of experience and a 700+ team spanning across multiple geographies - Netcore is revolutionizing the way marketers and product teams engage with consumers, through its AI-powered growth marketing platform designed to help marketers listen, analyze, and build 1:1 smart engagement.

Seeking a Smarter Way to Grow

Netcore's clients use its solutions to plan, execute, and monitor marketing campaigns across different channels such as Email, SMS, App, WhatsApp, and so on. Given limited budgets, the key ROI challenge for clients is to target the right customers, at the right time, on the right channels, and with the right message to maximize response rates and conversions. To counter these problems, Netcore created Raman—an innovative AI platform that analyzes huge datasets of historical and recent customer behavior to deliver smarter customer segmentation, improved targeting,

and sophisticated predictive modeling.

As clients rapidly adopted Raman alongside Netcore's core growth marketing platform Smartech, the performance of Netcore's real-time analytics modules began to decrease. The company needed a database able to handle write- and read-intensive workloads in parallel without any lag or drop-in efficiency. Its existing implementations of MySQL and MongoDB were unable to handle this workload efficiently, leading to slower model refresh and analysis.

Karthigai Muthu, Data Architect at Netcore Solutions, says, "We couldn't perform the required aggregation using MongoDB, and the delay in exporting the data for training our models and then pushing it back into the database increased the execution time. Refreshing the Send Time Optimization models for some of our larger clients was taking around 12 to 16 hours."

Debapriya Das, Head - Machine Learning at Netcore Solutions, adds, "AI systems are data-hungry, and slow access to data on MongoDB meant we could only build relatively unsophisticated models."

Optimizing Performance with Vertica

The Netcore Machine Learning team ran a proof of concept for Vertica Analytics Platform and quickly realized that the solution's ability to read huge volumes of data at high speed would solve its performance issues and accelerate the development of new AI solutions. The team started with a single Vertica resource pool on AWS and later added a dedicated three-node cluster for its AI engine, Raman. To protect data and enable faster recovery in the event of unplanned downtime, Netcore uses Vertica in Eon Mode, which stores a copy of all data in a shared S3 storage environment.



At a Glance

- Industry**
 Professional Services
- Location**
 Mumbai, India
- Challenge**
 Boost performance and scalability of AI platform to help clients maximize marketing campaign responses, conversions, and marketing ROI..
- Products and Services**
 Vertica Analytics Platform
- Critical Success Factors**
 - + 11x boost in query performance
 - + 10-20x reduction in weekly model refresh time
 - + Saves time and AWS resources with internal analytics and efficient Python integration
 - + Cuts time to market for new AI-powered solutions

"The performance boost from Vertica has been exceptional, with many jobs running more than ten-times faster."

Debapriya Das
 Head - Machine Learning
 Netcore Solutions

“Vertica UDx optimizes the AWS data transfer time and costs, as well as being easier and more efficient for our ML Engineers.”

Karthigai Muthu
Data Architect
Netcore Solutions

Contact us at:
www.vertica.com

Like what you read? Share it.



“The performance boost from Vertica has been exceptional, with many jobs running more than ten-times faster,” says Debapriya Das. “And because we no longer need to pre-aggregate the data, we can build deeper and more sophisticated models that empower our clients to reflect the real world more closely. One of our latest innovations—Predictive Segments—is a very complex feature, which looks at each user’s past web and app activities and predicts their propensity to purchase a product, to uninstall the app, to engage with a particular campaign, and so on. Vertica’s flexibility, built-in analytics modules, and performance helped us build a very time- and cost-efficient solution for our clients. We are planning to move some of our previous machine-learning work like Preferred Channel and other Send Time Optimization modules to Vertica as well, seeing the performance we can get out of it.”

Time-Saving Features

Machine learning requires ongoing data transformation, model training, and validation, which creates hundreds of temporary tables and large volumes of intermediate data. Using Vertica’s External Tables feature, Netcore can offload these temporary data sets, which keeps licensing costs low. Karthigai Muthu says, “Without External Tables, we would need another 6 TB of capacity in Vertica. The decentralized architecture also helps us to minimize AWS system resources, further optimizing our recurring costs.”

Using Vertica User Defined Extensions(UDx) has helped the Netcore Machine Learning team to run data analysis and processing scripts without the need to pull data out of Vertica first. Scripts written in Python can be registered as libraries in Vertica Analytics Platform.

Debapriya Das says, “Since we migrated the low-level data analysis scripts to the same instances where Vertica is running, this colocation made the overall process faster.”

Vertica further boosts performance by automatically parallelizing the script execution across multiple nodes. The team also uses Vertica’s built-in analytics functions wherever required, thus saving additional coding time. These features make RAMAN modules more time-efficient, which was not possible with Netcore’s legacy technology stack.

“Vertica UDx optimizes the AWS data transfer time and costs, as well as being easier and more efficient for our developers,” says Karthigai Muthu. “We can also use the machine learning libraries provided by Vertica; for example, we can predict using TensorFlow and PyTorch models on data residing in Vertica for predictive analytics.”

For model training, Netcore previously had to convert table columns into multidimensional arrays in system memory—a slow and resource-intensive process. With Vertica’s Array data type, Netcore can directly create multidimensional arrays, saving time and system resources.

Faster, More Intelligent Marketing

Not only did migrating to Vertica increase query performance, but it also cut the time required to refresh the various AI models. Different Raman modules for multiple clients can now be retrained or refreshed within a very short time, paving the way for deeper analysis.

Karthigai Muthu says, “Our Churn Management model has become really fast since we migrated it to Vertica. For clients, where end-to-end data preparation and model building - used to take nearly a day, we can now complete that work within 10 to 15 minutes. This is amazing as it saves us a lot of time, our team has already started working on model enhancements leading to better predictions.”

Netcore clients who combine Smartech with Raman not only get faster segmentation—helping them identify the right audiences for promoting their brands ahead of their competitors—but also AI-driven automation and

suggestions on how to optimize the content, timing, and strategy of campaigns.

Raman provides an entire suite of AI applications offering solutions around Predictive, Diagnostic, and Prescriptive analytics. It monitors data 24x7 to identify if something is going wrong or if a marketing KPI is taking a hit, without marketers having to worry about it, and even identifies the probable reasons.

For instance: Raman is fully capable of identifying if weekly revenues have started dropping at the earliest possible stage, and or if revenue from an app has started to fall because of high uninstall rates. Once the marketer validates and is satisfied with Raman’s observation, they can use Churn Management to understand which customers are likely to uninstall and then target them using a combination of Send Time Optimization, Preferred Channel, and Predictive Segments modules to halt the trend of decreasing revenue. In short, Raman is always learning..

Internally at Netcore, adopting the Vertica Analytics Platform has reduced AWS costs, increased efficiency, and accelerated the creation and training of new machine-learning models. Thanks to Vertica, the company can now bring new functions to market faster, which helps its clients maintain their competitive edge.

Kalpit Jain, CEO of Netcore Solutions, says, “Netcore has always enabled clients to achieve a better return on investment for their marketing budgets. With our new AI-powered offerings within Raman, helped by Vertica technology, we are making marketing automation faster, smarter, and more sophisticated. Ultimately, this helps our clients to deliver relevant offers and promotions to their customers, boosting both profitability and customer satisfaction.”