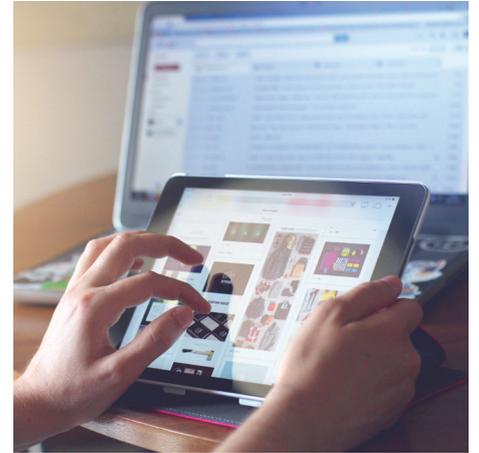


Propeller Ads

Introducing Vertica allows Propeller Ads to continually improve the performance of its advertising network.



Overview

Propeller Ads, a multichannel advertising platform, was founded in 2011, to create an optimal match between advertisers and content publishers. Advertisers are offered a wide range of modern ad formats over web and mobile, leveraging advanced marketing performance analytics to optimise marketing spend for over 100 million users in 195 countries. Every day, the company handles about 3 billion ad requests and generates 2 terabytes of new data.

Challenge

Propeller Ads offers a flexible and transparent market matching for publishers and advertisers, based on intelligent algorithms which ensures a fair market price for both parties. Advertisers can access the best targeted sites for their ads to increase their marketing ROI, ensuring the advertising campaign covers the widest audience and maximise conversion

rates for click-throughs. Publishers can identify the most profitable advertising campaigns for their media to receive maximum revenue with minimum effort.

The scale of the network, the data volume, and the analytics response time impose very challenging technical requirements to Propeller Ads' analytical platform. Their analytic engine must process in real-time huge volumes of transactions to optimise ad delivery and calculate marketing costs and performance metrics on the fly. In addition, advanced machine learning algorithms requires high performance processing to match the response time expectations of Propeller Ads customers.

The company started to suffer the rising cost of scaling and performance tuning of the existing MySQL-based data warehouse. Therefore, a decision was made to move to a more efficient and performant solution.

Solution

Propeller Ads started to test Vertica capabilities by using the free Vertica Community Edition, which was used to implement a few prototype data marts to boost the speed and agility of end-customer reports without changing the underlying enterprise data warehouse. Jacob Becker, Vice President of Propeller Ads,

“Advertisers and Publishers receive detailed statistics that reflect different viewpoints as quickly as possible.”

JACOB BECKER

Vice President
Propeller Ads



PropellerAds

At a Glance

Industry

Business Services

Location

Russia

Challenge

Requirement to analyse real-time advertising transactions and high-performance data processing for machine learning to optimise advertising campaigns.

Products and Services

Vertica Analytics Platform

Results

- + Increased speed of report generation
- + Increased number of metrics available for analysis
- + Provides a deeper understanding of the advertising process
- + Continuous improvement of the advertising rotation process

explains: “The solution demonstrated significant advantages in processing speed. Request execution time reduced ten-fold and some requests were executed 100 times faster compared to the previous technology”.

However, when the company decided to modernise and upscale the entire data warehouse, Vertica was not automatically selected, but compared against other candidates, including Hadoop and MongoDB. “Our main requirements were speed, scalability, and ease of use and maintenance through a small and agile engineering team,” explains Sergey Nikolenko, Head of BI Development at Propeller Ads.

Another important selection criterion was the ability to work natively with other key product of the business intelligence ecosystem. Since Tableau was already used to generate reports for Propeller Ads customers and internal users, the new platform had to provide efficient integration and support for advanced visualisation without introducing additional tools and software. Equally important was an efficient connection with existing ETL systems to load data into the new repository, while maintaining know-how and work from the previous implementation.



Propeller Ads launched a proof-of-concept using all analytical engine candidates, but Vertica surpassed competitors in all key benchmarks. Easydata, a key local Micro Focus® partner and Vertica system integrator, supported Propeller Ads for the full Vertica deployment and the solution administration. During the migration, Vertica specialists worked alongside Easydata to provide value-add consultancy during the migration and the performance optimisation. As a result, the data warehouse modernisation was a smooth process, lasting only six months.

However, a growing need for advanced analytics demanded not only a better performance on the existing architecture, but also to fully exploit new Vertica capabilities. The next step was to introduce a lambda architecture for advanced analytics based on Vertica. Lambda architecture is increasingly deployed by customer requiring high performance data processing, because it consolidates batch and

online data processing into a single data storage unit, accelerating data processing across different tasks and analytical workloads. The lambda aggregated data layer is designed to extract specific metrics from the data, analyse them, and generate reports in real or near real-time. Raw and unstructured data is collected as is, without changing formats, and stored for further processing.

A significant drawback of many lambda implementations, though, is that they are designed with very heterogeneous systems. The multiple integrations of heterogeneous components can often significantly affect the processing speed benefits for which the architecture was selected in the first place. Many lambda architecture implementations in the market do not allow SQL queries or easy integration with traditional BI systems. This makes it difficult for BI analysts to work and integrate with BI solutions.

Propeller Ads implemented a lambda architecture entirely on Vertica, and this helped to avoid such disadvantages. All components are leveraging the same storage. The aggregated data layer generates reports for advertisers and publishers, giving them dashboards with up-to-date, detailed statistics to analyse the effectiveness of ad impressions and the value of advertising campaigns. Since Vertica supports SQL query syntax and seamlessly integrates with BI systems, the new advanced analytics are not changing the existing processes of BI analysts.

Report generation is now carried out in near real-time mode, with no more than a minute delay for the most complex applications. Sergey Nikolenko comments: "Vertica storage capacity enables us to provide much more metrics to customers. Before Vertica, we collected only a few aggregated metrics, now any important analysis parameter can be stored in

the repository and be easily available to generate new metrics very easily. Thanks to the new architecture and Vertica, our advertising solutions now deliver approximately one hundred metrics, a ten-fold increase from before."

Propeller Ads also leverages Vertica FlexZone technology to ingest semi-structured data in the same data storage, thus realising a very effective Data Lake of up to 90 Terabytes over a single technology. Vertica FlexZone allows you to ingest into Vertica semi-structured data in its original format, while accessing it with SQL queries. The ability to work on both structured and unstructured data in the same way as a traditional DBMS is one of the key differentiators that convinced Propeller Ads to consolidate the data lake over a single technology.

"The storage capacity has increased hundreds of times compared to the previous architecture, where only aggregated data (e.g., averages



over a certain period of time) was maintained. Now, every ad show and every end user click is saved in our platform, which is very important to increase the accuracy of our server algorithms," explains Sergey Nikolenko.



A team of data scientists developed machine learning algorithms to optimise the mediation between advertisers and publishers. The ultimate goal is to improve advertising rotation, i.e., allow each marketer to test and identify the most profitable marketing mix. For example, Propeller Ads offers different payment models to advertisers including ad conversion user action. In this model, the advertiser pays for each new subscription or every download of a mobile app. Conversion depends on a number of factors, including what site the user came from, what time the interaction took place, what sites he visited before, and so on

Machine learning algorithms help determine the optimum format of each advertising campaign and increase conversions to maximise profits for all ad network members. The company data scientists also use the TensorFlow machine learning library, which is easily integrated with Vertica as a core repository.

“The storage capacity has increased hundreds of times compared to the previous architecture. Now, every ad show and every user click is saved in Vertica, which is very important for the smooth running and the accuracy of the advertising server intelligent algorithms.”

SERGEY NIKOLENKO
Head of BI Department
Propeller Ads

Contact us at:
www.vertica.com



the internet provider, in what city and country it is looking to advertise, mobile platform version, type of connection, i.e., 3G or Wi-Fi, etc. Our clients have access to comprehensive reporting on the services they receive.”

The power of Vertica’s processing speed and stored data volume opens up great opportunities for the company. Already, data that has not been processed before, for example, visitor’s browser history, is being collected. In addition,

the company can supplement this by purchasing data on users’ activities from third-party resources. Future plans include platform integration with other advertising server systems, which will allow Propeller Ads to further increase data volumes and the scope of collected information. “The ultimate goal is to provide an in-depth understanding of current processes so that we can continually improve our advertising rotation,” concludes Jacob Becker.

Results

Vertica provides Propeller Ads with a high performance, reliable, and scalable framework for the implementation of the lambda architecture, through which the company improves its advertising network every day, increasing customer satisfaction. “Advertisers and publishers receive detailed statistics that reflect different viewpoints,” Sergey Nikolenko notes. “Now we can collect and analyse the most diverse information about the user: geodata, the name of

