

REAL-TIME BIDDING

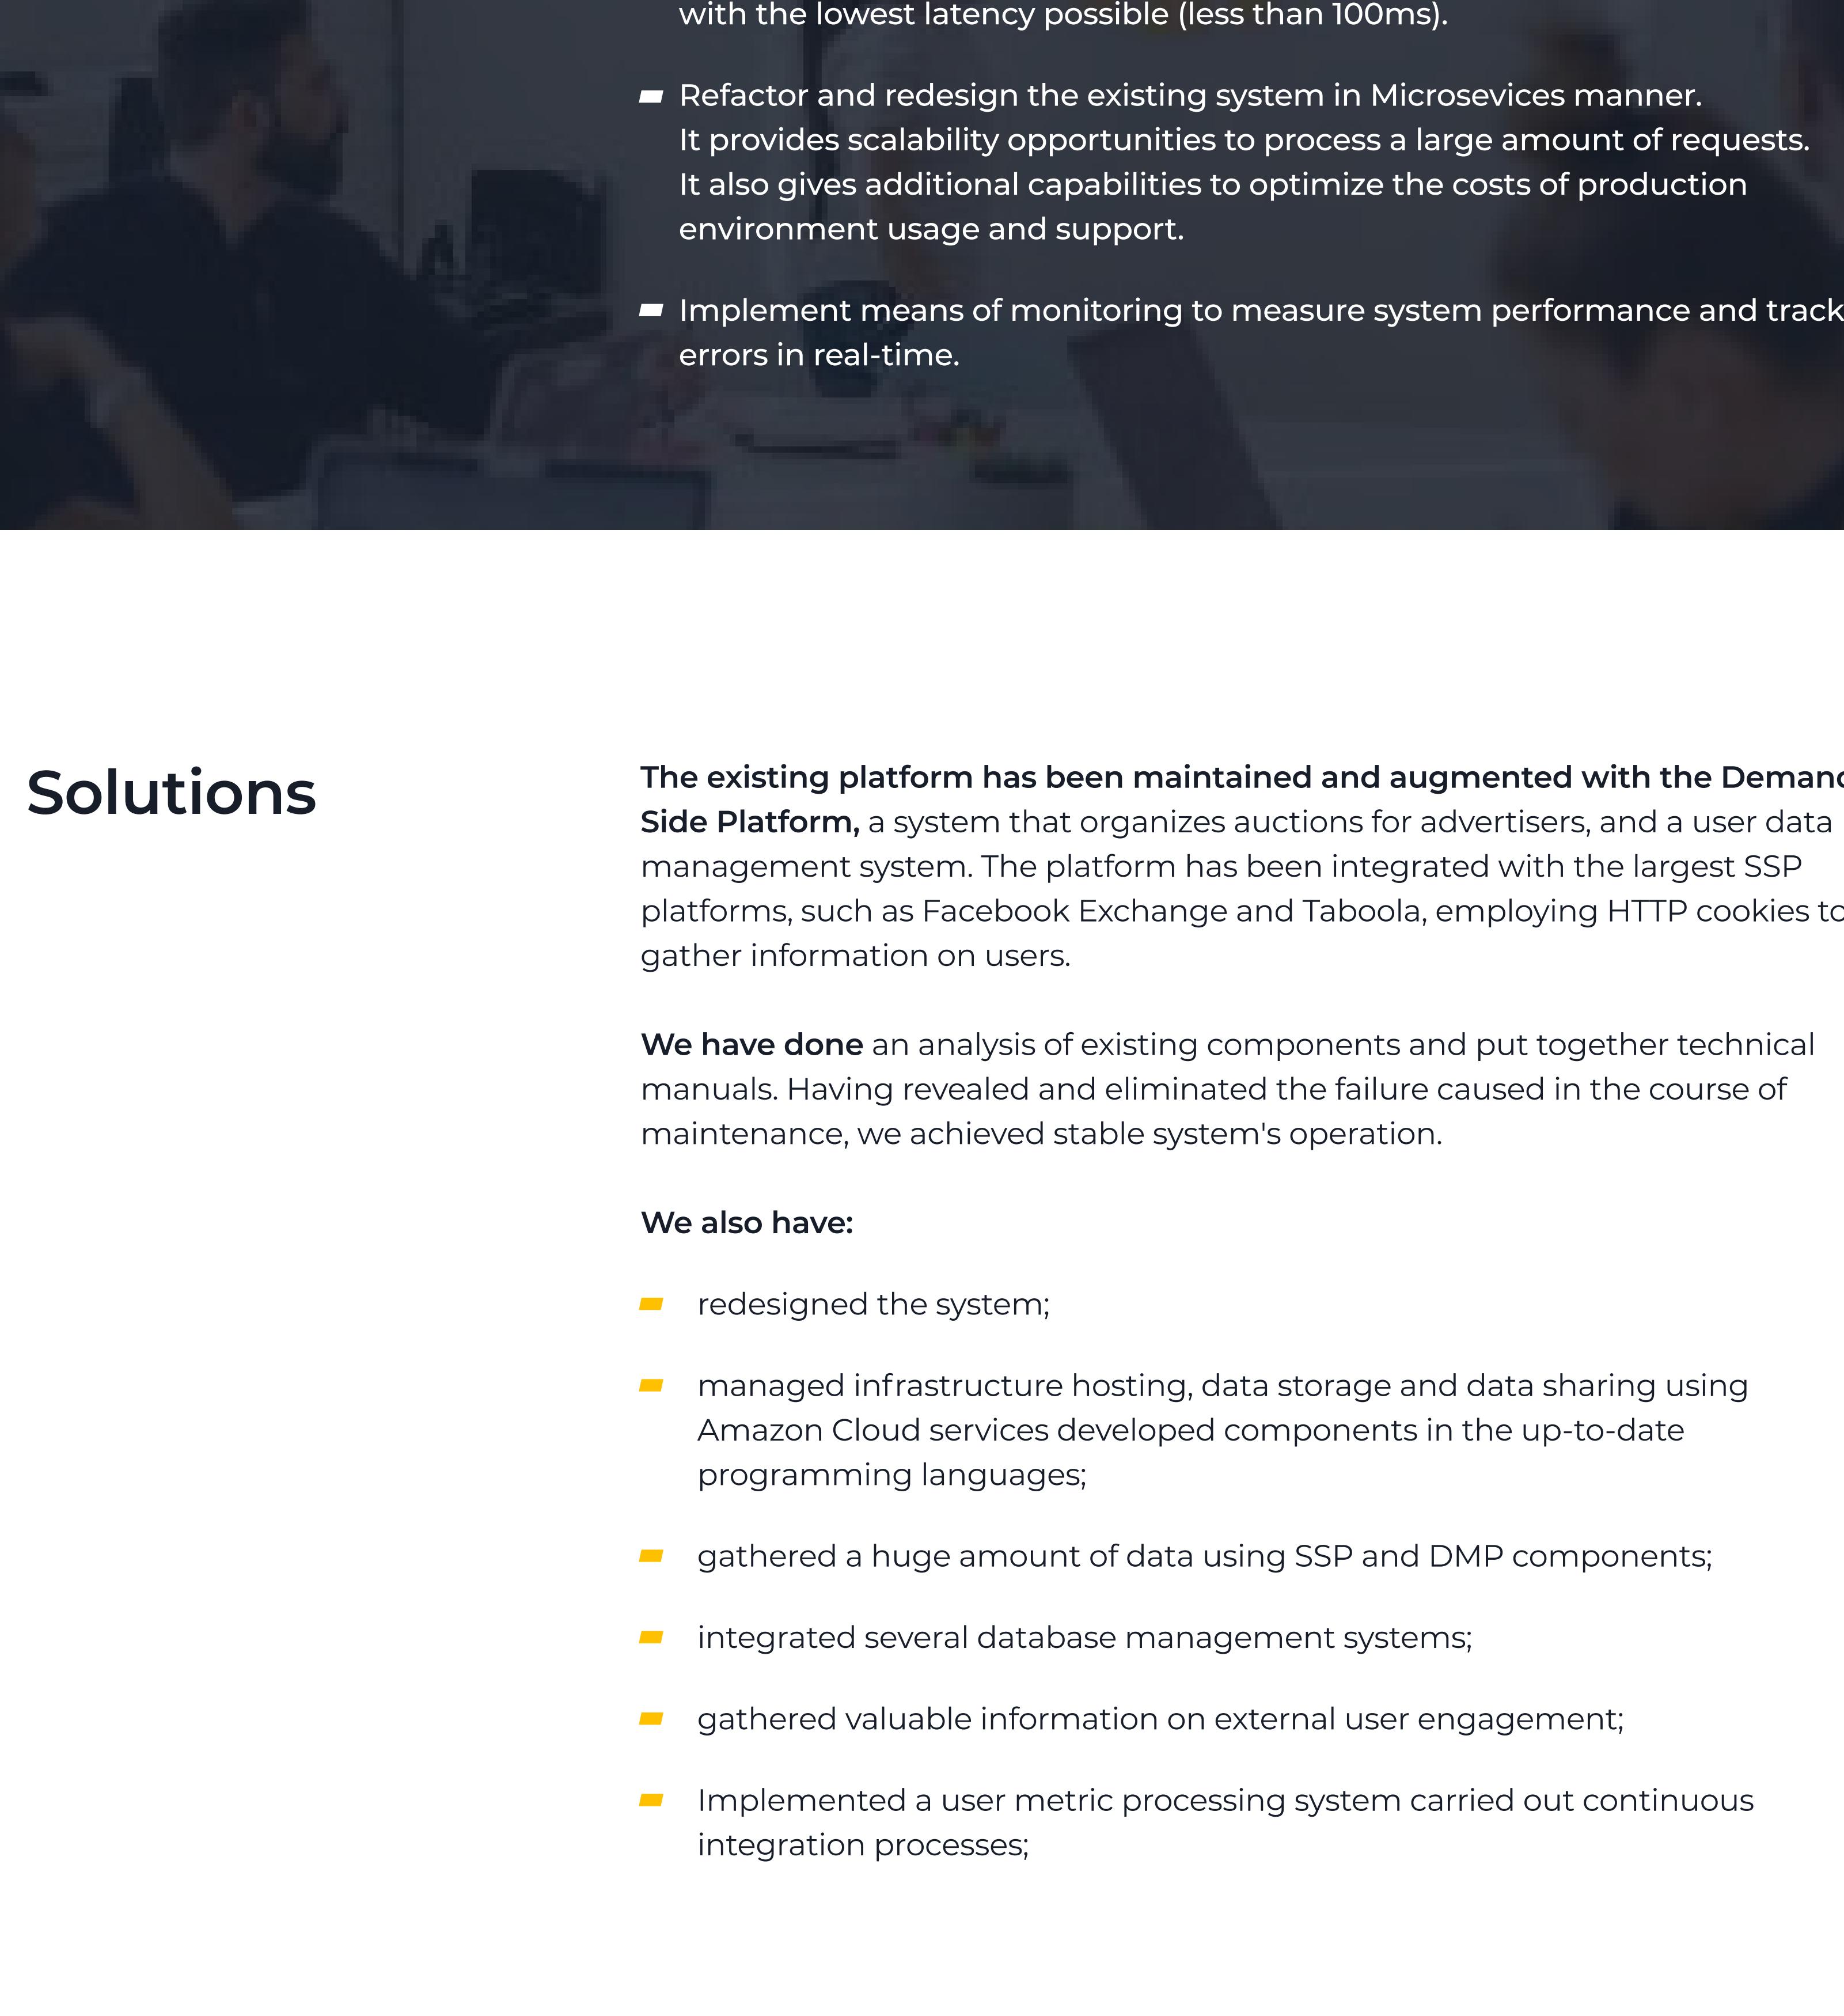
Real-time Bidding is a means by which advertising inventory is bought and sold on a per-impression basis, via programmatic instantaneous auction, similar to financial markets. With real-time bidding, advertising buyers bid on an impression and, if the bid is won, the buyer's ad is instantly displayed on the publisher's site.

About the Project

*instinctools company was involved in the Real Time Bidding project that uses the Personal Retargeting technique. It is about showing ads to users, who have visited an advertiser's site, however, failed to perform the conversion action. Personal Retargeting brings them back and shows each user a specific scenario, depending on their previous action: viewing an item, comparing an item to other modifications etc.

DSP (Demand Side Platform) makes each 1 click cheaper. While using Trading Desk, advertisers adjust bids according to the value of their target audience, as well as to the desired impressions and clicks volume. The platform compares the data with the information on traders and bids, buying impressions as cheap as possible and selecting the best matching audience. The whole process takes a split second: while the page is still loading.

Real-time Bidding



Challenges

- Support existing system functionality with lack of documentation and original developers' knowledge.
- Provide high-availability system to process a large amount of requests coming from SSP and advertiser sites. Each request should be executed with the lowest latency possible (less than 100ms).
- Refactor and redesign the existing system in Microservices manner. It provides scalability opportunities to process a large amount of requests. It also gives additional capabilities to optimize the costs of production environment usage and support.
- Implement means of monitoring to measure system performance and track errors in real-time.

Solutions

The existing platform has been maintained and augmented with the Demand Side Platform, a system that organizes auctions for advertisers, and a user data management system. The platform has been integrated with the largest SSP platforms, such as Facebook Exchange and Taboola, employing HTTP cookies to gather information on users.

We have done an analysis of existing components and put together technical manuals. Having revealed and eliminated the failure caused in the course of maintenance, we achieved stable system's operation.

We also have:

- redesigned the system;
- managed infrastructure hosting, data storage and data sharing using Amazon Cloud services developed components in the up-to-date programming languages;
- gathered a huge amount of data using SSP and DMP components;
- integrated several database management systems;
- gathered valuable information on external user engagement;
- Implemented a user metric processing system carried out continuous integration processes;

Technologies

The system design was based on well-known Amazon cloud technologies like Elastic Compute Cloud, Elastic Load Balancer, Simple Storage Cloud, Elastic MapReduce, DynamoDB, etc.

Aggregated reporting included processing of huge data amounts received from SSP and DMP components. It was implemented using Hadoop jobs executed in Amazon EMR. The results were aggregated in Amazon Redshift databases.

As a part of asynchronous operation processing, the several messaging platforms were used. NSQ provided distributed queues for high amount data processing. Sidekiq was used to implement business logic messages processing through the set of components written in Ruby.

Application metrics were processed using Graphite and were shown by Grafana.

Continuous integration processes were implemented based on Gitflow by using Github + Travis services.

For the needs of infrastructure hosting, data storage, data sharing the following services were used: Amazon EC2, Amazon ELB, Amazon S3.

The components were developed in modern programming languages. Go language capabilities gave an opportunity to implement performance critical items. Scripting functionality which included component integration, configuration and deployment were completely implemented with Ruby and related technologies.

Multiple database management systems were used based on the required functionality: DynamoDB and Aerospike used to collect huge amount of data related to outside user activity, Redis databases satisfied the needs of temporary data keeping, PostgreSQL and Amazon Redshift provided strong consistency storage and processing of data related business activity and reporting.

Rollbar service was used to track errors occurred in a production environment.

Key Features

The client has obtained a stable version of the platform together with the build products, as well as an option to solve integration issues quickly and at a low cost in the future.

The augmented platform benefits are:

- Tracking the system in a real-time mode
- Integrating with other systems within its own framework
- Performing the expected functions fast and accurately

The client was pleased with our fulfillment of the requirements and the recommendations of advertising partners. We are currently continuing to maintain the stable operation of the platform.

Technologies



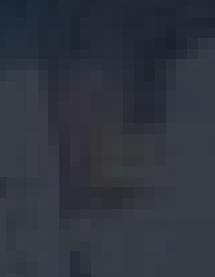
Go language



NSQ



DynamoDB



Aerospike



Redis

Sidekiq

PostgreSQL

GitHub

Rollbar

Gitflow

Ruby

Grafana

Hadoop jobs

Travis

Amazon Web Services:

- Elastic Compute Cloud
- Elastic Load Balancer
- Simple Storage Cloud
- Elastic MapReduce
- DynamoDB
- Amazon EC2
- Amazon ELB
- Amazon Redshift
- Amazon EMR
- Amazon S3