

Customer Case Study

Timeful

Customer Case Study

Timeful

Benefits

- Improved key metrics monitoring by processing the entire production data set instead of sampling subsets
- More effective data-driven product design in a much shorter cycle
- Redirect one FTE in data analytics to focus on problem solving instead of data analysis for other teams

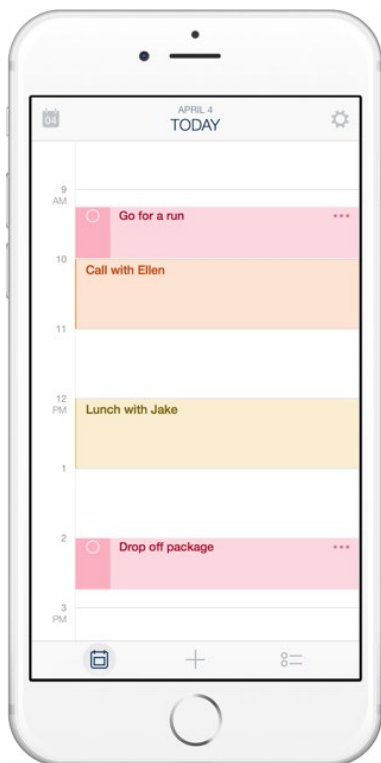


Summary

- Timeful needed to improve their product design and recommendation algorithm by broadening access of production data to non-engineering personnel and hourly monitoring of key metrics
- Their existing solution of multiple Postgres databases was too slow to support hourly monitoring of the entire data set, and too cumbersome to be used by the broader team
- Timeful used Databricks Cloud to implement hourly monitoring and provide non-engineering teams access to production data, significantly improving recommendation quality and product design cycle

Business Background

Timeful’s mission is to help people manage their time better: track commitments, fit items on TODO lists in their busy schedules, and ultimately find the time to develop good lifestyle habits - such as exercising regularly. Many people have difficulties with time management because traditional calendars force them to manage time manually, by painstakingly searching through a cluttered series of appointments and reminders for a sliver of opening. Timeful augments the traditional calendar with data science in a smartphone application to reduce the manual work required by the end user taking all the stress and complexity of manual planning away. The Timeful application uses machine learning to recommend a personalized schedule for each end-user based on availability, preferences, and behavior.



To help people accomplish items on their TODO lists, the Timeful algorithm allows one to enter deadlines with each item, and recommends a time for each item before the deadline is reached. The recommended schedule improves as one provides feedback by rescheduling the TODO (rejecting the recommendation), or checking TODOs off the list (accepting the recommendations).

Similarly, Timeful uses recommendations to help people develop good lifestyle habits. To find time for a specific habit, one sets the frequency (e.g., 3-times per week) and Timeful will then recommend times to undertake that habit. Similar to TODOs, the recommendation quality also improves as one reschedules recommendations and marks habits as “done”.

Challenge

To create a positive user experience, the Timeful team relies heavily on data analytics to improve product design and monitor the personalized schedule recommendations. Both of these data analytics tasks proved to be challenging in Timeful's environment because Timeful stores its data in multiple Postgres databases. Having multiple data sources without a centralized and intuitive platform for data analytics created barriers for the designers and product managers to access the data they needed to inform the product design process. As a result, they often had to call on data analysts and scientists to help them with specific analysis tasks, severely limiting the effectiveness of both teams.

The performance of Postgres also prevented Timeful from monitoring the quality of their recommendation algorithm effectively. The application generates terabytes of data from its recommendations, which must be monitored continuously to ensure the satisfaction of the end-users. However, the Postgres databases were too slow to support complex and frequent monitoring across the entire data set. As a result, monitoring was performed by sampling a small subset of the data.

Solution

Timeful chose Databricks Cloud as a centralized, high-performance data processing platform to make up for the limitations imposed by multiple Postgres databases and achieve the following capabilities:

- **Continuous monitoring of recommendation quality:** Due to the speed of Apache Spark, Databricks Cloud is able to monitor the entire production data set on an hourly basis without sampling. Deploying the monitoring programs at scale was also simple with Databricks Cloud — Timeful simply had to use the “Jobs” feature to schedule the computation of quality metrics at the desired intervals. The Jobs feature also allowed Timeful to send automated alerts to the engineering team in the event of an error condition. This gave Timeful the ability to automatically monitor production algorithms without human intervention.

- **Easy access and exploration for non-engineers:** The interactive workspace feature in Databricks Cloud enabled designers and product managers to rapidly inspect data and explore hidden relationships in a fast, visual, and iterative way. The designers and product managers are also empowered to perform analysis by themselves without the technical expertise from data analysts or scientists because the interactive workspace was intuitive and easy to use.

To get started with Databricks Cloud, Timeful was able to easily deploy it in its Virtual Private Cloud (VPC) of Amazon Web Services (AWS) within days, and import data into it from production data sources through Amazon S3 with Databricks Cloud's built-in APIs. The zero-management capability of Databricks Cloud also meant that Timeful could achieve this without investing in DevOps personnel or sacrificing the productivity of its engineers.

Benefits

Through Databricks Cloud, Timeful gained significant benefits by improving its production quality monitoring and product design while avoiding expensive infrastructure and DevOps personnel:

- **More accurate recommendation quality monitoring:** Databricks Cloud's performance and speed enabled Timeful to perform complex monitoring over its entire dataset on an hourly schedule automatically, instead of using only sampled data, eliminating the possibility of undetected bad recommendations. Additionally, Databricks Cloud was also able to automatically send out alerts if the monitored metrics failed to meet Timeful's standard - further ensuring the quality of recommendations.

“The speed of Databricks Cloud and the power of the Spark are unparalleled. Post implementation, we’ve been able to run complex monitoring over our entire dataset on an hourly basis in an automated manner. The value of that simple automation for my team alone is worth the investment. For the first time, we don’t feel like we’re three steps behind with a fast and comprehensive monitoring system.”

– Gloria Lau
VP of Data, Timeful

- **More effective data-driven product design in a much shorter cycle:** Designers and product managers are able access production data via Databricks Cloud. As a result, they were able to understand Timeful users better by running in-depth investigations ranging from funnel analysis to cohort studies.
- **Focus data analysts and scientists on problem solving rather than accessing data on behalf of others:** Teams can access the data independently without needing data analysts and scientists to act as the liaison. This enabled data scientists to focus more on problem solving, instead of servicing the rest of the company on their data needs. As a result, Timeful was able to shift a data analyst into a front-end development role.

“Before switching to Databricks Cloud, our processes to access production data were too slow and the Postgres databases were too cumbersome for our non-engineering teams. Databricks Cloud allowed them to focus on creative problem solving and attend to the most important needs of all—those of our users.”

– Gloria Lau
VP of Data, Timeful

Evaluate Databricks Cloud with a trial account now.
databricks.com/registration