tecton + tide

Tide Saves Time and Money by Automating Data-Driven Business Processes With Real-Time Machine Learning

> *"Tecton makes it simple to leverage this idea of having a 360-view of your customer and then feeding that into your various operational applications... What is the value that we generated? For credit specifically, we leveraged a roughly 50% increase in approval rate, while decreasing the losses by about roughly 5%, and fraud transaction monitoring was even more extreme, so we 4x the chance of fraud for every blocked transaction while blocking 20% fewer transactions. Overall, we caught more fraud, but also we did this by blocking fewer transactions."*
>
> — Hendrik Brackmann, VP Data, Tide

OVERVIEW

About:

Tide (Tide Platform Limited) is a UK financial technology company providing mobile-first banking services for small and medium-sized businesses. Its customers can set up a current account and get instant access to various financial services, including automated bookkeeping and integrated invoicing. Established in 2015, Tide is one of the first digital-only finance platforms in the UK to provide current accounts for businesses. Tide is headquartered in London, with offices in Sofia, Bulgaria, and Hyderabad, India.

Challenge:

To address real-time use cases and to scale their operational ML efforts across the board, Tide began developing an internal feature store. Not only did they encounter hurdles incorporating both batch and streaming data into their feature store, but they also quickly realized that building their basic feature store would take many months and would require 3 additional full-time employees just to maintain it.

Solution:

With support from Tecton's engineering team, Tide was able to roll out Tecton's fully integrated feature platform into production in just 6 weeks.

Results:

With Tecton, **Tide has not only cut the time it takes to deploy a model from 2–4 months to just 1 month, they have also improved model accuracy with 7x more features and now deploy 2x more models than they had previously.** With Tecton, Tide has designed and deployed real-time data products like risk analytics or fraud detection in record time.

Learn more:

Watch this talk where Hendrik Brackmann shares Tide's journey from developing their own feature store to adopting Tecton's feature platform.

Challenge: Limited functionalities of an in-house feature platform

To build ML-based products like real-time fraud detection and risk assessments to automate new account approvals or ML-driven transaction-to-invoice matching, Tide honed in on developing the ability to make predictions online in real-time. To address these real-time use cases and to scale their operational ML efforts across the board, Tide began developing an internal feature store. They were able to start serving some of their features online but encountered hurdles incorporating both batch and streaming data into their feature store. Furthermore, although their in-house feature store meant no vendor lock-in and gave them complete control of the roadmap, it presented the following challenges:

1. **Adding features was hard**
    - Adding features from new topics required significant engineering overhead
    - No support for adding features from Snowflake
    - No support for stream aggregations

2. **Changes to production were unpredictable**
    - Expected changes from training did not hold up in production environments
    - Approval rates were lower than anticipated during training
    - Had to do feature engineering twice for backtesting

3. **Slow iteration times**
    - Adding new features required extensive engineering support
    - Strain on hiring; difficult to attract top talent
    - Over 50% of time was spent on creating feature pipelines

Finally, Tide estimated that completing a basic feature store to meet their requirements would take 6–9 additional months, and they would need to hire 3 full-time employees just to maintain it.

Solution: Overall cost efficiency of an enterprise feature platform

Considering the financial and time costs of building their own solution, Tide decided to evaluate Tecton's enterprise feature platform. With support from Tecton's engineering team, Tide was able to roll out Tecton's fully integrated feature platform into production in just 6 weeks.

> *"Buy versus build? While the features themselves are strategically core to us, getting a good understanding of our customers and building good features that are actually predictive is core to what we want to do strategically, but the computation orchestration layer is very much not. So we choose Tecton for optimized orchestration."*
>
> — Hendrik Brackmann, VP Data, Tide

As part of Tide's unified ML data layer incorporating both historical and real-time data, Tecton's feature platform enabled Tide to accomplish the following:

Easy re-use of high-quality features. Tide's data science and engineering teams heavily re-use features to power different models. For instance, Tide's credit risk and fraud detection services re-use many features, like those related to transaction history or those related to not fulfilling government obligations. Transaction fraud models perform better with features computed from short-term transactions whereas credit models perform better when they leverage longer time periods. Using Tecton's opinionated feature build framework, Tide is now able to fine-tune the features for any use case in a fraction of the time it would take them to build new features from scratch.

Ease of backtesting. Before deploying changes to production, Tide is now able to estimate the holistic impact of these changes on production environments. Teams can now regularly run experiments to measure how changes to rules or ML impact overall system results, approval rates, or historical risk realizations. These pre-deployment estimations closely match production performance.

Improved deployment time. Tide uses Tecton to re-use pre-built feature types, greatly reducing workloads previously involved when adding features. With the introduction of Tecton as a core component to their uniform data layer, Tide reduced its overall delivery time from model conception to deployment in production by 50%. By using 7x as many features on average as they did before, Tide's model accuracy is also greatly improved.

Results: A feature platform to streamline the development and deployment of real-time ML use cases

With Tecton, Tide has not only cut the time it takes to deploy a model from 2–4 months to just 1 month, they also have improved model accuracy with 7x more features and now deploy 2x more models than they had previously. They've also alleviated pressure on engineering hiring and project management and made high-quality feature engineering core to the company's ML-driven product strategy to automate as many processes as possible.

With the help of Tecton's feature platform for real-time ML, Tide has successfully created a set of tools and solutions designed to improve customer experience.

> *"With Tecton, it's very easy to renew high-quality features, specifically within the same domain. Backtesting is easier and is also actually correct. We also saw a significant improvement in deployment times. The overall delivery time has been significantly reduced when it comes to building additional features and machine learning models."*
>
> — Hendrik Brackmann, VP Data, Tide

Tide and Tecton, in practice

An example of how Tide uses Tecton in practice is its automated risk analytics data product. To ensure Tide customers can safely open business current accounts straight from their mobile phone by scanning a photo ID, Tide must be able to approve new clients at the time of their request. To minimize friction for customer onboarding, reduce the number of hours spent on manual reviews, and speed up new account approvals, Tide set out to automate the process with immediate, real-time predictions of credit risk per customer.

#1: Risk Analytics: Automated Credit Risk Assessment for New Account Approvals

Challenge: Batch and Manual Credit Risk Detection

SMBs apply for business accounts or loans and expect a decision from Tide within seconds. During this new customer approval and onboarding process, Tide must simultaneously evaluate risk while providing a smooth customer experience. Tide must tread carefully when approving such customer requests because, for example, if a loan is approved but later defaults, Tide accrues financial losses. On the other hand, rejecting loans for qualified applicants is a bad customer experience, decreases conversion rates, and increases acquisition costs.

Tide's risk evaluation must be highly accurate to avoid future loan defaults and must provide results in under 200 milliseconds so that customers do not experience delays during the sign-up process. For this to happen, Tide faced the challenge of serving features online based on streaming data— their previous models used only batch data and offline features.

Solution: Real-Time Risk Assessment

Tide built models to detect different types of risk across the hundreds of thousands of transactions that occur daily on their platform. These models leverage fresh features from transactional information, behavioral data, third-party data (e.g., credit bureau scores), and streaming data to accurately assess credit risk for individual accounts in real-time. Tide automated the use of high-quality features like transaction history or features related to not fulfilling government obligations to compute risk scores for individual customers. They found that their credit models were more impacted by transaction history from longer time periods than fraud models that benefited more from the short-term transaction history.

Result: $600k Cost Savings

By implementing high-quality features in their credit risk assessment models, Tide has increased approval rates by 50% and decreased the average loss per credit underwritten by 5%. Furthermore, the company has reduced spending on manual reviews for new accounts by over $600K/year.

#2: Real-time Transaction Fraud Detection

Another example of a finance tool Tide built and deployed with Tecton is real-time fraud detection which thousands of SMBs depend on to safely process transactions.

Challenge: Failure to detect fraudulent transactions leads to poor customer experiences

To build customer trust and create a safer user environment, Tide needed to improve a fraud detection solution that could continuously profile and identify risk across user sessions at the exact moment behavior potentially indicated fraud without increasing false positives or negatively impacting user experience.

Solution: Online inference powered by fresh features from real-time data

Tide uses transactional data, behavioral data, and third-party data (e.g., credit bureau scores) to learn from past fraudulent behavior in order to predict, in seconds, the likelihood that a new transaction is fraudulent. As part of Tide's unified machine learning data layer, Tecton allows teams to orchestrate and leverage historical and real-time features for both training and backtesting, as well as powering models with real-time and batch source feature pipelines for real-time inference.

To build their fraud detection system, Tide automated the use of high-quality features like transaction history or features related to not fulfilling government obligations. They found that transaction fraud models benefit more from short-term transaction history than credit models, which are more impacted by transaction history over longer time periods.

Results: Tide significantly reduced false positives and improved real-time fraud detection

By implementing high-quality features in their fraud risk assessment models, Tide has decreased blocked transactions by 20% all the while seeing a 4x increase in fraud likelihood for every blocked transaction.

> *"We have events coming on Kafka as well as batch jobs that are feeding into Tecton. For us, Tecton is a unified layer of historical and real-time features. If you are looking for a particular data set, Tecton makes it available, both historically in order to do analytic tasks on top of constrained models, estimate rule impact, and then also makes it available in real-time, so that your systems can use it in order to make low latency decisions."*
>
> — Hendrik Brackmann, VP Data, Tide