



How QuanticMind Moved to Continuous Delivery and Eliminated Deployment Delays Using Split





Overview

QuanticMind was struggling to increase delivery of new functionality for their B2B customers. The team wanted to aggressively increase their release cadence— from once every 3 months to a continuous delivery model— but they saw risk in rapidly releasing new features across their diverse and growing customer base.

In addition, a cumbersome branching strategy meant changes took precious time to propagate into production.

Split helped QuanticMind separate code deployment from feature release, letting them press the gas pedal on release cadence while gaining full control over the customer exposure to new functionality.

Data points

- Reduced their release cadence from 3 months down to a high-velocity continuous delivery model
- Eliminated a 1 week delay of new feature rollouts with an improved branching strategy
- Avoided 3 months of engineering effort to improve an in-house feature flagging system

“ Split has transformed the way that we deliver software. Granular targeting of features to customers has made us more responsive, and data from Split gives us an audit trail in case any issues arise ” – A Back End Tech Lead at QuanticMind

Situation

As a rapidly growing software company, QuanticMind started hitting some challenges to their feature rollout strategy. Customer growth had been rapid, and a growing number of enterprise customers started using the platform. QuanticMind felt it was becoming increasingly important for them to phase the rollout of new features to reduce the risk across their broad customer base.

At the same time, the team was looking at the next step of upgrades to their internally-built feature flagging system. They estimated it would take the engineering team 3 months to build new capabilities such as audit flag changes or establish permissions to scale up the number of people who could make changes to feature flags.

Strategically, QuanticMind wanted to adopt a continuous delivery model to meet the demands of the business. The company had already accelerated from a 3 month release cadence to 1 week, but even the weekly schedule wasn't enough to support their rapidly growing business. Merging branches of code into the trunk was a big issue that hindered release velocity, often causing rewrites of existing code to accommodate a new feature. This meant that when sales teams promised new features, there would be a frustrating delay in custom availability.

Choosing Split

The engineering team at QuanticMind evaluated options to upgrade their in-house feature flagging system. They needed to meet the demands of continuously releasing code into production and make it easier to turn on features in production.

As the team took a look at open source options, they learned that many of the available tools were designed for just one part of the development stack. QuanticMind needed something that could support the full-stack, both the frontend JavaScript and backend Java/Scala components.

The off the shelf options left a significant amount of work for the engineering team to integrate into the codebase. Many were designed to only handle the rollout of features by individual users. They were not able to handle rolling out features by groups such as a team or specific account— a key requirement for QuanticMind.

QuanticMind found out about Split and saw how easily the Split Feature Experimentation platform could be integrated into their codebase. Split also provides an easy to use UI to manage granular feature rollouts by customer account, supports scalable permissions, and tracks change history for regulatory compliance needs.

An added benefit of Split was a pre-built integration with Slack. Split's Slack integration ensures team members are instantly aware of feature flag changes. The messages in Slack also provide color-coded detail on the state of a flag to make them easy to read.

Getting started with Split was simple for the team - they were able to do much of the initial setup work themselves, leveraging Split's expert services as needed. Overall, the process only took a couple of days.

According to David Samuelson, Back End Tech Lead, QuanticMind, "Before Split, our process for turning 'on' or 'off' a feature in production was for one of our engineers to make a change to a config file. And that feature state change would be global to all customers." This approach had several problems. The first issue was that an engineer had to be distracted every time a salesperson wanted to show a customer a new feature. It didn't take much time but was more of a nuisance."

Result and Future Plans

QuanticMind started using Split in select areas of their platform. The Split footprint within the company has grown significantly since and is now being used across the entire engineering team. Anyone can add feature flags, and teams can coordinate across the organization on things like feature naming and tagging to improve standardization.

Every new feature is now put behind a feature flag, making every feature rollout an experiment to ensure production stability. With Split, QuanticMind can safely phase in new functionality across their customer base.

They have also significantly improved their branching strategy to reduce risk of production failure. Branching off the main trunk of code is minimized with small amounts of code being incrementally merged into the trunk and kept dark from customers until tested and ready. QuanticMind is able to actually test features in production since the "blast radius" of a feature in a production can be minimized.

The engineering team works hand-in-hand with the Product Management and Quality teams in deciding what to put behind a flag and determining the granularity of what to break up behind flags. Over time, the team also cleans up flags that are no longer used. Split provides an easy way to see if a version of a flag has no traffic being routed to it, effectively meaning it can be killed or deleted.

Overall, QuanticMind has decreased their release cycle times and increased the pace of feature delivery, while minimizing risks to the customer experience. Split has played a key role in helping QuanticMind continue to accelerate innovation and eliminate deployment delays.

About QuanticMind

QuanticMind, the Platform for Smarter Marketing, is the industry's most intelligent performance marketing cloud for data-first marketers. Leveraging insights from proprietary machine learning intelligence, the QuanticMind platform--anchored by campaign optimization, predictive analytics, and powerful reporting capabilities--unifies the full power of all customer and marketing data to deliver peak campaign performance. Based in Silicon Valley with 11 global locations, the QuanticMind platform powers the world's top consumer brands to increase pipeline, drive revenue, and improve overall marketing ROI. For more information, please visit QuanticMind.com.

About Split

Split is the leading platform for feature experimentation, empowering businesses of all sizes to make smarter product decisions. Companies like Salesforce, Vevo, and Twilio, rely on Split to securely release new features, target them to customers, and measure the impact of features on their customer experience metrics. Founded in 2015, Split's team comes from some of the most innovative enterprises in Silicon Valley, including Google, LinkedIn, Salesforce and Databricks. Split is based in Redwood City, California and backed by Lightspeed Venture Partners, Accel Partners, and Harmony Partners. To learn more about Split, contact hello@split.io, or get started for free at www.split.io/signup

