



Vida Health Case Study

INCREASING PRODUCT DELIVERY WHILE LOWERING RISK

Vida Health is a digital health company offering a full suite of personalized health programs and one-on-one coaching and therapy.

Vida's horizontal platform integrates leading behavioral science and data science to drive health outcomes for members, including those with chronic physical and behavioral health conditions. Vida leverages Split's feature flagging and deep experimentation capabilities to build and deliver innovation faster with less risk.

THE RESULTS

- Reliably increased its mobile application release velocity from one per month to one per week by using Split for phased rollouts and testing in production.
- Realized a 30% increase in the number of features that make it into each release train now that all new code is behind a feature flag and tests are run in a production environment.
- Product experiments used to take a week to set up and analyze. Now they are run instantaneously so data scientists are free to focus their expertise on building new therapeutic models for their customers.

THE SITUATION

Vida wanted to implement phased rollouts for their mobile application to meet its rapid delivery goals while lowering the risk of production failure. Their engineers created an in-house feature flag solution that leveraged a database for turning features on or off. When a major feature was ready for a phased rollout, a column would be added to the database with a Boolean on/off indication.

The Postgres database stored user information, which required an engineer to go to an admin page and manually turn the feature on for each user. Once they were ready to roll out to 100% of their user base, a SQL query could be used to turn the feature on for all users. This method gave them the fine-grained control they desired, but there was no randomization applied to the users and the process became burdensome.

When a flagging implementation is database-backed, an engineer has to make a patch for every feature release. Any code change is inherently risky and

comes with the potential for error, plus this is a constant drain on engineering resources. As its user base and business requirements grew, Vida engineers knew they needed a more scalable solution.

USING SPLIT

The Vida engineering team was under pressure to support a more rapid release cadence for their mobile application, however, they lacked confidence in the slow and potentially error-prone manual system they had in place. They brought Split in for a proof of concept because they liked the idea of a platform that supported both feature flagging and experimentation.

Vida's first use of Split was to evaluate a new machine-learning model for ranking their information retrieval system vs. using older business logic. Once they were comfortable with Split for back-end analysis, they started creating splits for minor UI enhancements such as testing new animations or testing functionality of different end-user devices. Because the features were behind a flag, they were able to easily roll back when issues were uncovered.

With this realization, momentum grew within the engineering team. They were ready to apply Split to upcoming major application features. Some of the new features incorporated many new elements that impacted one-fourth of the application, so the team felt they needed the reliability and predictability that Split offers.

Vida is confident in test results because they know Split runs the same algorithm every time. Previously, every test process was started from scratch by a different engineer, which led to inconsistent metric definitions. 'Do these results include outliers this time or did we remove them?' 'What's the standard process?' Now they don't have to worry about tracking previous algorithms because Split ensures consistency and standardization from test to test.

“As we ramped up our usage of Split, engineering teams could see everything went smoothly. This increased our confidence enough so that now every feature and code change we make is wrapped in a Split feature flag.”

– **Russell Melick, Engineering Manager, Vida Health**

ALL CODE CHANGES ARE BEHIND A FEATURE FLAG

Today’s standard process is to have every feature and code change behind a feature flag for both the front-end mobile application and their back-end Python Django legacy web application. Now that every code change is behind a feature flag, engineering and product teams collaborate on the ramp plan. They use Split’s integration with Slack so everyone can see when splits get ramped.

ITERATE AND IMPROVE ON FEATURE DEFINITION

Vida’s engineering and product teams use Split as a tool to iterate and improve features. If the minimum viable feature doesn’t test well with users, rather than kill it right away they will work to make it better, then execute another user test. The rapid feedback loop helps engineers refine functionality and has become a standard process they account for in roadmap planning.

TEST IN PRODUCTION

The flexibility provided by feature flagging allows Vida’s quality engineering team to test against live production environments. In the past, when engineers built a new version of the application they deployed only to employees. They couldn’t release test code to the application store, which meant waiting until the version was ‘completely perfect’ before releasing.

Engineers would point to a staging environment to test the new build, but they weren’t always confident that the errors they saw were from actual code or from some configuration error in the staging environment. With Split, they upload a version to the application store that contains new code and bug fixes and expose them only to Vida engineers for testing prior to customer release.

Test accounts are created to validate the onboarding flow of new functionality and are quickly cleaned up when testing is done. They also test against pre-existing accounts for their weekly releases to ensure ample amounts of test data are available.

Vida engineers still leverage a staging environment when working on code or for back-end functional testing, but they no longer have to manage separate code branches for the mobile application. Monitoring code stability in production and phasing rollouts gave Vida engineers the confidence to meet aggressive release cycle time goals.

HOW AWS SUPPORTS SPLIT’S SOLUTION FOR VIDA HEALTH

Split can scale to serve feature flags and experimentation easily by leveraging an architecture built on AWS. Split runs all its services in containers with Kubernetes on AWS and EC2, allowing it to scale rapidly as needed to support all the rollout and experimentations of customers like Vida Health. And to faster serve all the necessary data for the flags and their rollout definitions, Split caches all its data in AWS Redis in front of its Atlas MongoDB clusters.

To support all experimentations, Split also leverages Amazon S3 to scale to whatever volume of data is needed, and it uses EC2 with Databricks to do the metric calculations. Metric calculation results are stored in Aurora to then calculate all the statistics on demand. This architecture has enabled Split to effectively set up an experimentation platform that can scale to all of Vida Health’s needs. Finally, AWS also allows Split to build a highly available system for Vida Health leveraging the multiple availability zones and regions that AWS supports.

CHOOSING SPLIT

For Vida, having one platform that combines both feature flagging and experimentation was a mandatory requirement. Split was more in line with its needs compared to other options, as Vida didn't want engineers working on experiments for web page design and didn't need dynamic HTML code generated. They needed a full-stack solution that would help engineering and product management experiment with both their mobile application and their back end-system to deliver a seamless customer experience. Split's straight forward user interface has helped product managers get involved with phased rollouts and experiments.

FUTURE USE

Vida engineers have worked with their data science team to get them comfortable with the analysis and results from Split. With Split, they can run their product experiments instantaneously, with no wait time for analysis. Vida plans to pull Split impression data into Google BigQuery where their data scientists conduct broader analysis for building new customer features and engineers train machine learning models. The inclusion of Split data in the Vida centralized data warehouse allows split impressions to be incorporated into a broader data pipeline managed and analyzed by the Vida data science team.

ABOUT SPLIT

The Split Feature Data Platform™ gives you the confidence to move fast without breaking things. Set up feature flags and safely deploy to production, controlling who sees which features and when. Connect every flag to contextual data, so you can know if your features are making things better or worse and act without hesitation. Effortlessly conduct feature experiments like A/B tests without slowing down. Whether you're looking to increase your releases, to decrease your MTTR, or to ignite your dev team without burning them out—Split is both a feature management platform and partnership to revolutionize the way the work gets done. Switch on a [free account](#) today or schedule a [demo](#) to learn more.



Switch It On at split.io

