

Split Intelligent Results Engine

Measure the Impact of Features on the Metrics that Matter

At a Glance

The Split Intelligent Results Engine is a metrics-first statistical engine built for engineering teams that handles data ingestion from any source, calculates metrics based on that data and joins the data to features, automatically highlighting significant changes due to feature releases across the metrics that matter most.

Benefits

- Easily experiment on every code change to see how it will impact product KPIs
- Ensure that everyone in your organization is making decisions based on the same industry standard statistical analysis
- Manage metrics globally, across any project in any language
- Choose which metrics to measure for an experiment at any time, for any experiment version
- No additional developer work required to associate events to visitor impressions

Capabilities

- **A/B/n Testing:** Granular test plans and automatic randomization and bucketing of visitors for experimentation
- **Data Ingestion and Metrics Calculation:** Bring in any event stream; Robust set of metric types, including sum, count and average
- **Automatic Attribution of Changes In Metrics to Experiments:** Split intelligently attributes changes in metrics to specific features
- **Statistics Engine:** Industry standard analysis of feature impact on metrics

Measure Feature Impact Against Product Metrics

As organizations grow their software products and digital offerings beyond an initial feature set, it becomes increasingly challenging to show how recent feature investments impact the business. Enterprises need access to real-time analytics on how customers are interacting with these features. And, these insights need to relate back to the full set of product KPIs.

It is not enough to simply report on website conversions. The impact of new features must be measured against low level metrics, such as exceptions or page load times. Metrics that measure the overall customer experience are critical as well, such as the number of actions or time spent in the application. And, ultimately feature changes should tie all the way back to the key business metrics, such as revenue or units purchased.

To have confidence in product decisions, all the conclusions from the data must be statistically sound and trusted.

The Split Difference

The Split Feature Experimentation Platform is the only solution that can measure actual features in your product against all your key product metrics. The Intelligent Results Engine module handles data ingestion from any source, intelligently pairing it with Split's own data stream on who saw what version of a feature. This means that product managers can use global, organization-wide metrics for any experiment, and do not need to specify a metric for an experiment ahead of time. Developers do not need to have knowledge of the event data associated with the code they are building, nor do they need to write code to link events to visitor (e.g. users, accounts, stores, whatever type of traffic is being tested) impressions.

Split lets you measure features against all the metrics the organization cares about, with flexibility and robust options for calculating metrics.

Importantly, Split does all this while maintaining data privacy - your sensitive data does not need to be shared with Split.

The Split Intelligent Results Engine works with the same core Split Feature Flag capability that is built to serve different versions of a feature to visitors based on any visitor attribute without ever sending attribute data to Split servers.



Examples

Experiment on every bit of code you release. Customers can use the power of Split to do things such as:

- Test the impact of a new feature on new sign-ups and revenue from a consumer application
- For a ride-sharing app, test the impact of a code change on app crashes or rides per user
- Evaluate if a new onboarding flow accelerates usage of advanced features in a B2B app
- Test if a change in processing of REST API requests for machine-to-machine communication improves data refresh time in a banking application
- Measure both the impact and number of visitor actions based on an algorithm change

Features

A/B/n Testing

Define feature experiments and create roll-out plans that execute different code in your software for different visitors. Split randomizes visitors and handles bucketing visitors into different cohorts.

Data Ingestion and Metrics Calculation

Ingest an unlimited amount of data from any source. Split collects the data, stores it, and builds metrics. Metrics are continually calculated for every treatment of every feature, and are automatically reset if any change is made to the experiment.

Automatic Feature and Event Attribution

Using the event timestamp and customer ID, Split will automatically attribute events to a matching experiment version, feature, and treatment experienced by a visitor.

Feature-specific Results Dashboards

Quickly see all the metrics that experienced a statistically significant change for a feature experiment.

Statistics Engine

Built-in industry standard statistical calculations. Split maintains a Type 1 Error Rate of 5%, a Type 2 Error rate of 20% and a Minimum Detectable Effect of 0.1 standard deviations (i.e. ~4 percentile lift). Split automatically displays results as they are calculated, along with the statistical significance of a result, allowing early results to be previewed.

How it Works

Split starts by collecting the impression data of every feature experiment. Impressions occur whenever a visitor is assigned a treatment (i.e. “variations”) of a feature. In parallel, Split ingests any external data stream via the Split Event API.

Once the data is processed, the stats engine begins analysis. Split automatically joins external events to visitor impressions. During this process, Split removes visitors from the test if the visitor was exposed to multiple treatments of the experiment. This could happen to a visitor that moves from New York to London, where geographic location is a variable in the test.

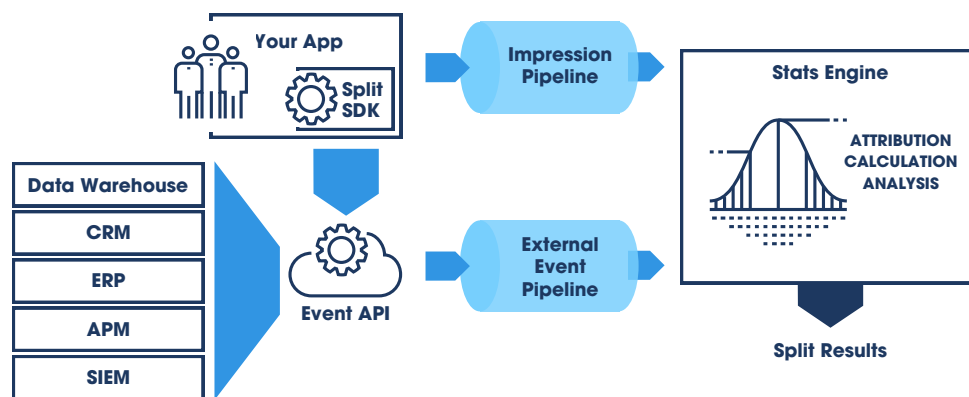


Figure 1: Split Intelligent Results Engine

Then, Split looks at each treatment of a feature, for each version of an experiment, and calculates all the metrics around the ingested external event data. If the experiment definition for a feature changes, Split automatically resets all metrics to ensure the integrity of the data.

Finally, the stats engine analyzes the data by comparing how the metric changed for those in one treatment versus a selected baseline.

With the Split Intelligent Results Engine, you can look at metrics for an experiment at any point in time, but Split will only color the change red or green if it is deemed to have a statistically significant negative or positive impact.

About Split

Split is the leading platform for feature experimentation, empowering businesses of all sizes make smarter product decisions. Companies like Vevo, Twilio, and LendingTree 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 Accel Partners and Lightspeed Venture Partners. To learn more about Split, contact hello@split.io, or start a 14-day free trial at www.split.io/signup.