

Rules Engine: Multiple Experiments and Deliveries per Flag



Status: Currently in Closed Beta. General Availability expected by May, 2023- only available in Accelerated and/or Scale tiers.

Accelerate your experimentation velocity

The biggest difference between Full Stack (legacy or simply FS) and Feature Experimentation (FX) is that all experiments, targeted deliveries, and Multi-armed Bandits (MABs) are based on feature flags. By eliminating the legacy A/B Tests in favor of Rules Engine and A/B Test rules, all A/B Tests are essentially Feature Tests (from legacy), and winning variations can be rolled out **without developer involvement**.

Today, all FX customers can create flags with up to one A/B Test rule and one Targeted Delivery rule. With this feature, we eliminate the cap on up to two rules per flag. This means that, with a single flag, customers can run several concurrent A/B Tests, MABs, and Targeted Deliveries.

This is a game changer for experimentation velocity: if FX customers follow the recommended workflow below, they can minimize developer involvement and enable non-technical collaborators to run many experiments and targeted deliveries **without the need to even tell a developer after the flag is implemented and tested**.

This feature has the potential to double or triple a customer's experimentation velocity alone. It also unlocks a new category of use cases for the industry: feature personalization, or the ability to tailor individual features to one or many audiences simultaneously.

Personalize your features

With Multiple Experiments and Deliveries per Flag, FX customers can follow the recommended workflow below to seamlessly run experiments, gain valuable insights, and roll out winning variants – all without needing a developer after the flag is implemented and tested. By enabling FX customers run more than one experiment and one targeted delivery at a time, we supercharge their ability to personalize features at scale by enabling them to release multiple versions of a feature to different audiences at the time same, all using the same flag.

How Rules Engine works

Rules Engine is unique to FX and enables customers to create a series of A/B Tests, MABs, and Targeted Deliveries on a single flag.

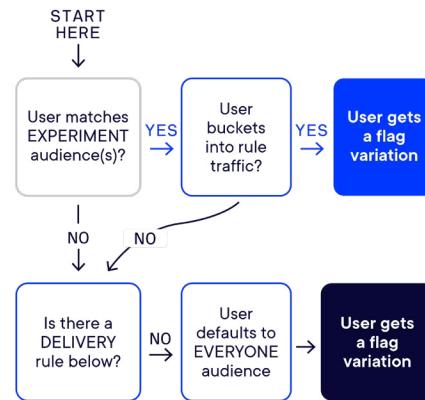
It works by sequentially evaluating each rule. While the evaluation logic is complex, it was created intentionally to maximize value from a customer's traffic by algorithmically bucketing their users in the first experiment or Targeted Delivery for which they qualify, instead of disqualifying them at the first opportunity.

Order of operations

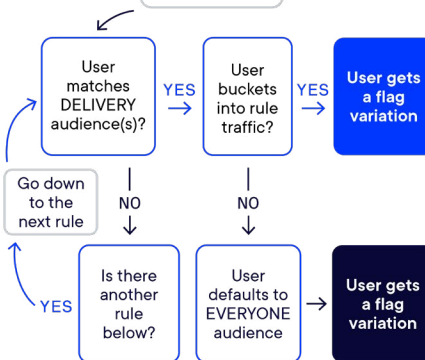
When a decision needs to be made for a user for a given flag:

Without Multiple Experiments per Flag

Experiment Bucketing




Targeted Delivery Bucketing

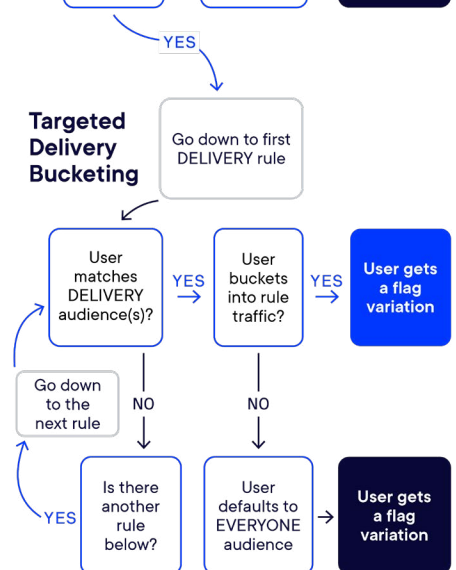


With Multiple Experiments per Flag

Experiment Bucketing




Targeted Delivery Bucketing



Recommended Workflow

FX customers can get the most value from Rules Engine and Multiple Experiments and Deliveries per Flag by following this workflow for all features (and flags):

1. A non-technical user or developer creates a Flag in the FX UI with variables for each part of the feature anyone may want to experiment on in the future. This can be thought of a feature or flag specification or spec.
2. The flag spec and corresponding code snippet from the FX UI can be put into a Jira (or workflow tool) story for a developer to implement.
3. A developer implements the flag following the flag spec, with sensible error handling (should there be an issue with the flag) and unit tests for the flag itself, including all its variables. Developers should test their features with multiple values for each flag variable, including empty or null values.
4. Once the developer is confident that the flag works as expected, they may deploy the code with the flag off. This is sometimes called a dark release.
5. No further developer involvement is required. Non-technical users may now add rules to the flag's Rules Engine pane, including A/B Tests, MABs, and/or Targeted Deliveries. They may turn the flag and rules on and experiment to their heart's content, without the need to align with the Dev team.
6. As A/B Tests are run and winning variants are identified, customers may pause or remove the A/B Test rules and add Targeted Delivery rules to roll out winning variations, **all without needing a developer.**



Q&A

Is Rules Engine available for Full Stack (legacy) projects?

No. This is only available in Feature Experimentation.

Is Multiple Experiments & Deliveries per Flag available for Full Stack (legacy) projects?

No. This feature requires Rules Engine and Rules Engine is only available in Feature Experimentation.

Why is there a different order of operations for evaluating experiments and targeted deliveries in Rules Engine?

The quirkiness is that Experiment rules (A/B Test & MAB) are not evaluated the same as Targeted Delivery rules. In both, if the user doesn't match audience criteria, the next rule is evaluated. For experiments, if the user is then not bucketed into the experiment's traffic allocation, the next experiment rule is evaluated. For targeted delivery, if the user matches the rule's audience but isn't bucketed into its traffic allocation, the rule engine jumps to the last rule (Then, for everyone else...) instead of evaluating the next rule.

At Optimizely, we're on a mission to help people unlock their digital potential. Our leading digital experience platform (DXP), equips teams with the tools and insights to create and optimize in new and exciting ways. Now, companies can operate with data-driven confidence to create hyper-personalized experiences. Building sophisticated solutions has never been simpler. Optimizely's 900+ partners and 1100+ employees in offices globally are proud to help over 9,000 brands, including Electrolux, Uber, Visa, WSJ, Santander, The Red Cross and Mazda enrich their customer lifetime value, increase revenue and grow their brands. Learn more at [Optimizely.com](https://www.optimizely.com)