

LaunchDarkly →

3 keys to maximize the business value of every product feature

Experiment your way to
continuous customer delight.

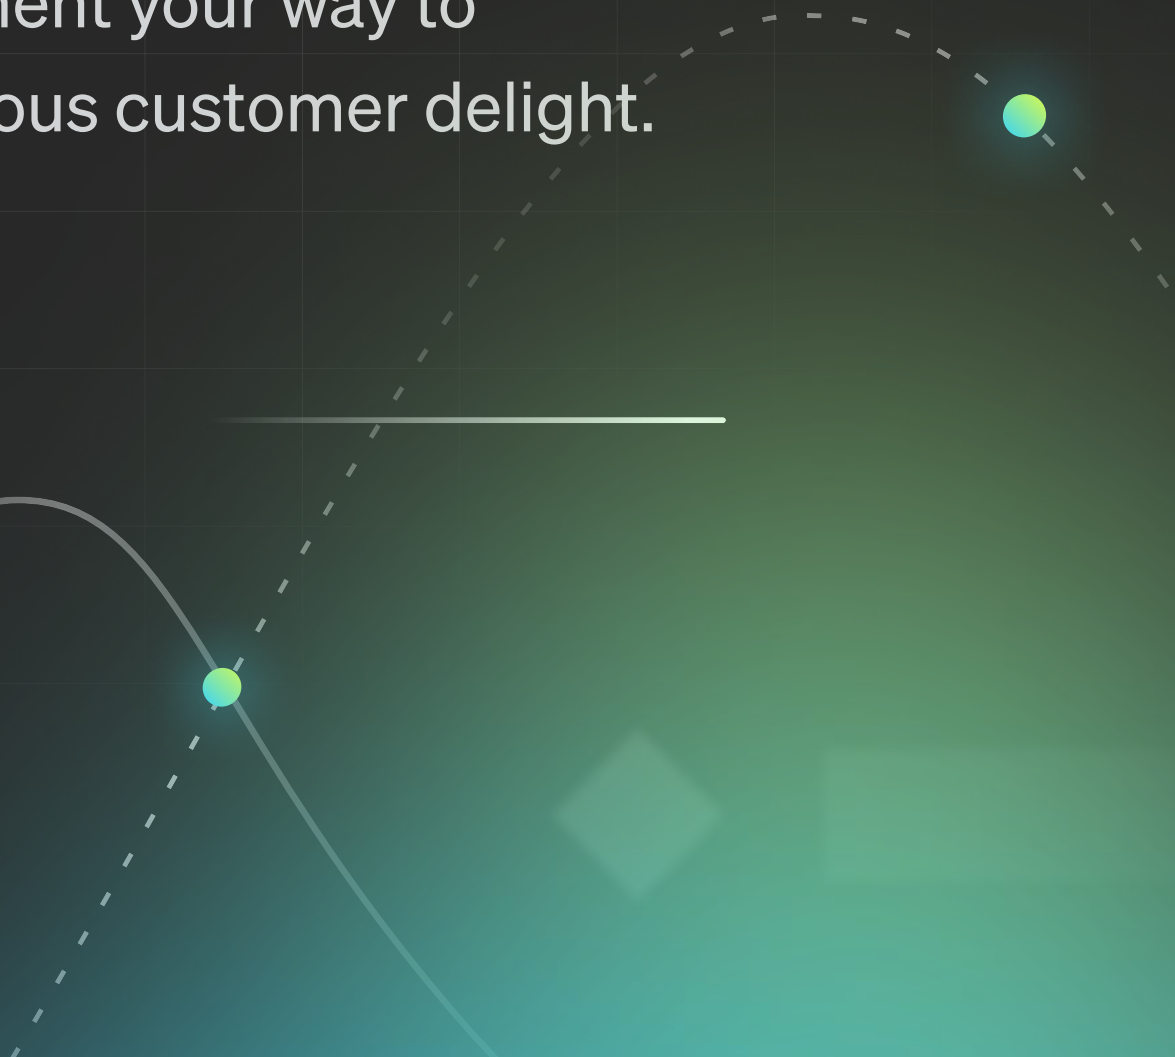
A decorative graphic consisting of a dashed white line that curves upwards from the bottom left towards the top right. Two glowing cyan circular dots are positioned along this line. The background features a dark green-to-teal gradient with a faint grid pattern and several blurred diamond shapes.

Table of Contents

Introduction	3
01 Top 3 barriers to impactful experimentation	4
Gut decisions + poor metrics = poor results	
A bad back-end spoils a good front-end	
Silos and unscalable processes hinder experimentation	
02 How LaunchDarkly can help	7
Improve the business performance of every feature	
Run back-end and infrastructure optimization experiments	
Accelerate safe innovation at scale	
03 Customer success stories	13
CCP Games	
Ritual	
04 Get started	18

The need to continuously ship features that drive business value

Leading companies like Expedia, Booking.com, Google, and Microsoft appear to have cracked the experimentation code. They run hundreds or thousands of tests a year to continuously improve the business value of their products. For example, in one account, Microsoft was generating hundreds of millions of dollars of additional revenue every year as a result of controlled experimentation.¹

Unfortunately, most organizations struggle to run even a few experiments that are truly valid. In our experience helping software teams enhance customer experiences, LaunchDarkly has observed three key challenges that block most organizations from successful experimentation. In this guide, we'll unpack these challenges and explain how LaunchDarkly enables you to overcome them.

LaunchDarkly helps organizations continuously measure and improve their digital products through experimentation. By shipping features that show proven impact, you drive predictable revenue. What's more, we enable customers to optimize every layer of the stack—from front-end to back-end to the edge and beyond. Finally, we help them achieve a clear line of sight across teams, products, and user experiences, bringing the big picture into focus and enabling safe innovation at scale.

In this guide, learn the secret to overcoming the three major barriers to software product experimentation.

¹ Fabijan, A., Dmitriev, P., Olsson, H. H., and Bosch, J., "The Benefits of Controlled Experimentation at Scale," 2017 43rd Euromicro Conference on Software Engineering and Advanced Applications (SEAA), Vienna, Austria, 2017, pp. 18-26, doi: 10.1109/SEAA.2017.47.

Top 3 barriers to impactful experimentation

//

Many firms run no more than a few dozen (experiments per year) that have little impact. If testing is so valuable, why don't companies do it more?²

Stefan Thomke

Professor, Harvard Business School

² Thomke, Stefan, "[Building a Culture of Experimentation](#)," Harvard Business Review, March-April 2020.

01

Gut decisions + poor metrics = poor results

Organizations often get stuck in the “just keep shipping” mindset of continuous delivery, leaving them little time or resources to determine what customers want or how best to give it to them. Without data to inform decisions, software teams are forced to go with their instincts. Unfortunately, when software development companies “trust their gut” about customer preferences, they can be wrong as often as 90% of the time.³

Plus, even when data is available, it frequently isn’t useful for driving business impact. Too often, teams settle for engineering metrics like lead time for changes and deployment frequency. While these measurements are helpful for meeting efficiency benchmarks, they have little-to-no effect on the customer experience. Decisions about product and feature changes should instead be informed by customer behavior analytics and collected feedback. But this data isn’t always available to developers and engineers.

02

A bad back-end spoils a good front-end

To offer the best customer experience, your application performance must be top-notch. This requires the ability to run A/B tests and other experiments on back-end infrastructure, components, and services. While solutions for cosmetic front-end experiments abound, few exist for the back-end.

³ Fabijan, Dmitriev, and et al., “The Benefits of Controlled Experimentation at Scale.”



As a result, developers, platform engineers, and SREs lack insight into how new software changes impact application reliability. What's more, they lack the data they need to fine-tune features with a view toward optimizing system performance. Frankly, it doesn't matter how good your front-end features are if the underlying back-end components are glitchy and unreliable. Until product delivery teams can run effective server-side experiments, their product experience will likely fall short of its potential.

03

Silos and unscalable processes hinder experimentation

At many organizations, the tools and processes for running a product experiment are disconnected from the tools and processes required to deploy the winning treatment of an experiment. That is, running experiments and deploying experiment results require separate implementations. This disconnect leads to fewer experiments, less engagement data, and a longer time-to-value.

Teams will often run small, discrete experiments without much trouble. But as they attempt to scale experiments, the work becomes more difficult, the process grows convoluted, and the communication pipeline starts to break down.

Lastly, experimentation tooling that fails to integrate with your tech stack causes yet more problems. You're forced to stitch data together across silos, deriving unreliable insights at a pace that's far too slow to meet customer demands—and with a process that's far too complex for non-technical users to leverage.

All of the above challenges put continuous experimentation and optimization out of reach.



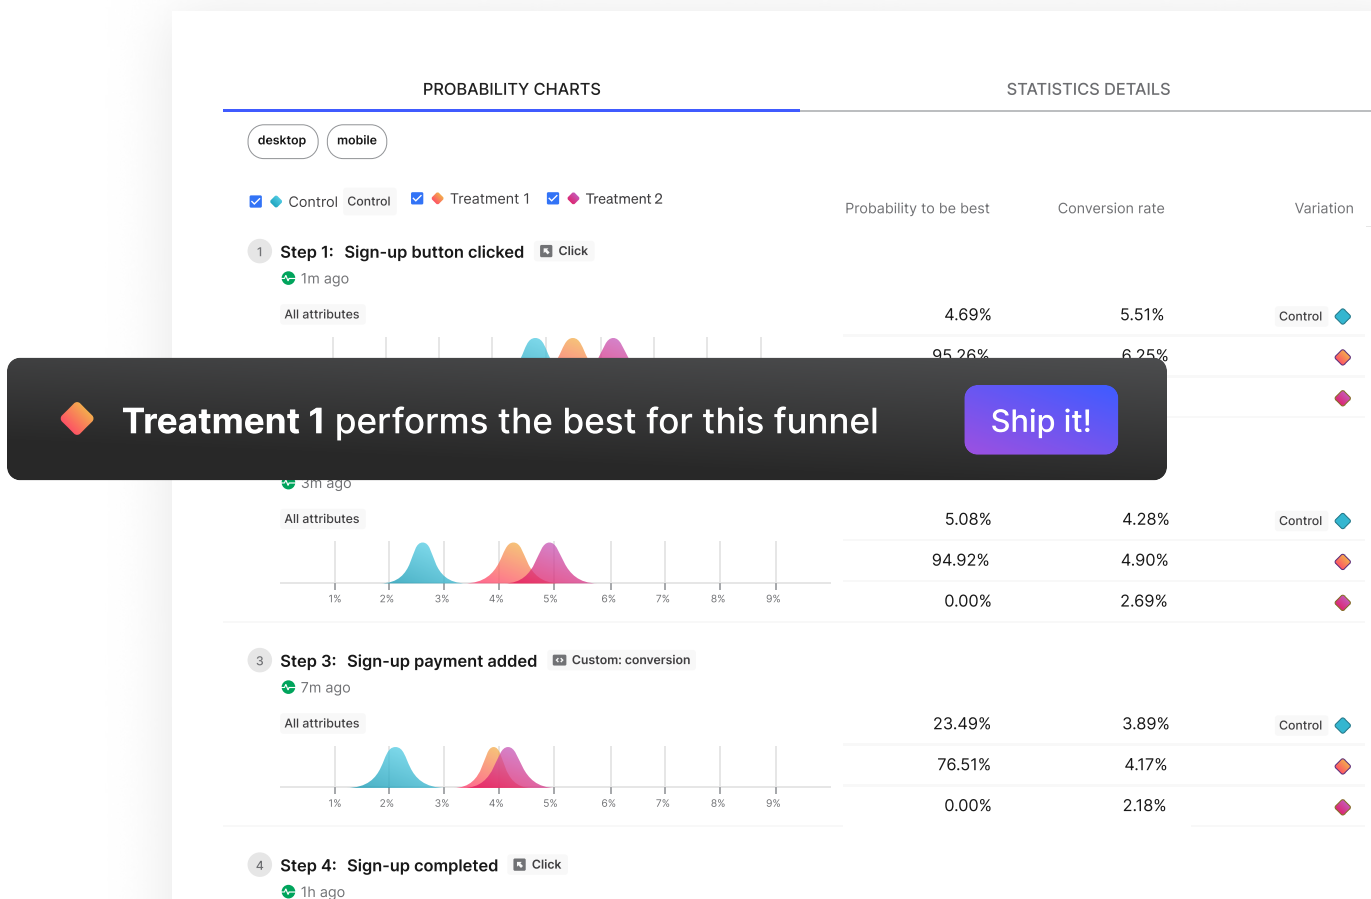
How LaunchDarkly can help

01

Improve the business performance of every feature

Rather than relying on guesswork or gut decisions, LaunchDarkly gives your teams the power to create metrics that align directly with business objectives —and easily integrate them into the feature lifecycle. From there, you can make smarter, data-driven decisions that optimize the value of every feature. And you can deliver more of what your customers want.

With LaunchDarkly’s out-of-the box funnel optimization [experiments](#), you can measure every product change, gaining the trustworthy data you need to boost key performance indicators (KPIs) across user engagement, acquisition, retention, and more. From there, you can go deeper, drilling into measurements on specific features that go far beyond basic events and clicks.



LaunchDarkly also lets you visualize the business impact of every product change—through metrics like conversion rates or traffic spikes—helping your teams clearly and confidently understand how product investments are driving progress. This becomes especially valuable for B2B companies that contend with limited data points, as LaunchDarkly's statistical engine generates reliable insights even with smaller sample sizes.

In environments where swift, informed decisions are imperative, LaunchDarkly becomes the linchpin that **empowers your teams to act quickly**.

02

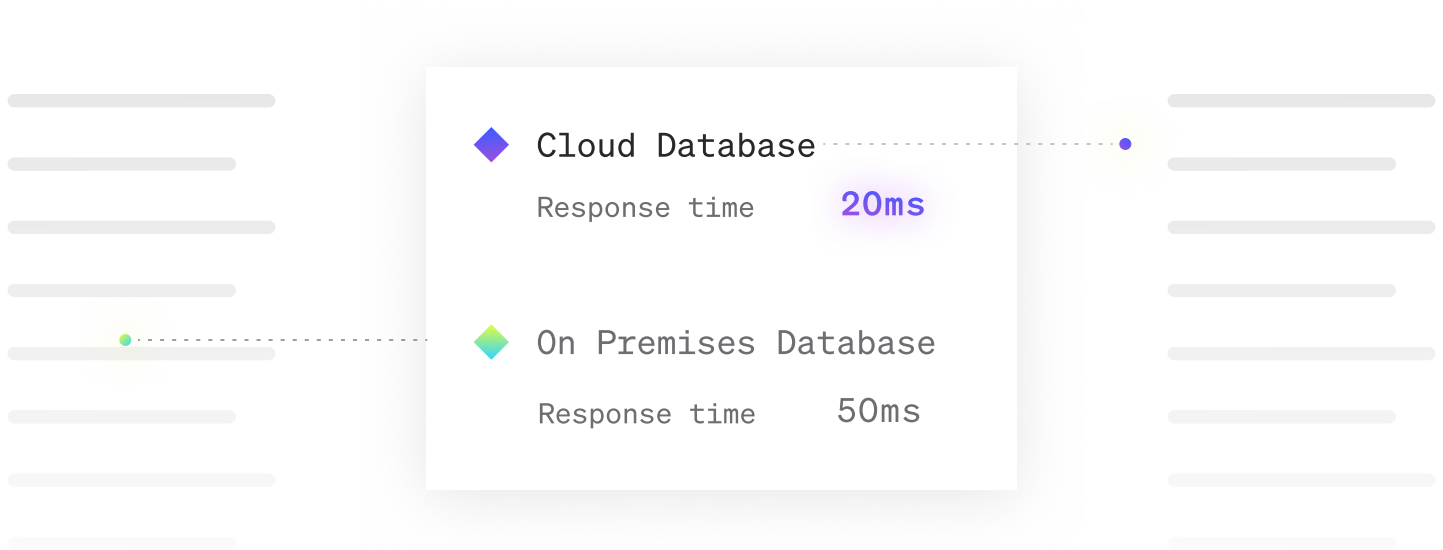
Run back-end and infrastructure optimization experiments

To compete in today's world, your applications and digital products must perform fast and reliably. With LaunchDarkly, you can optimize every layer of your software through targeted experiments. Moreover, you can achieve the back-end performance and availability to consistently deliver the exceptional experiences today's customers demand.

Leverage the wide breadth of LaunchDarkly's software development kits (SDKs) to experiment confidently across diverse environments, extending all the way to the edge. Whether you're testing new components or tweaking existing ones, the ability to optimize your full stack puts you in control.



For example, with LaunchDarkly, you can easily run an experiment to measure the impact of a new API on page load times. Or you can see how a new database implementation affects system performance. Or how a search feature in your mobile app impacts latency.



With LaunchDarkly, you can also quickly identify any bottlenecks or problems in production as you push changes to your application and infrastructure. Our integrations with popular [application performance monitoring \(APM\) solutions](#) like Datadog, Honeycomb, and New Relic provide clear, comprehensive visibility to all stakeholders. And they allow engineers and non-technical teams to remediate issues associated with an experiment by flipping a single toggle.




03

Accelerate safe innovation at scale

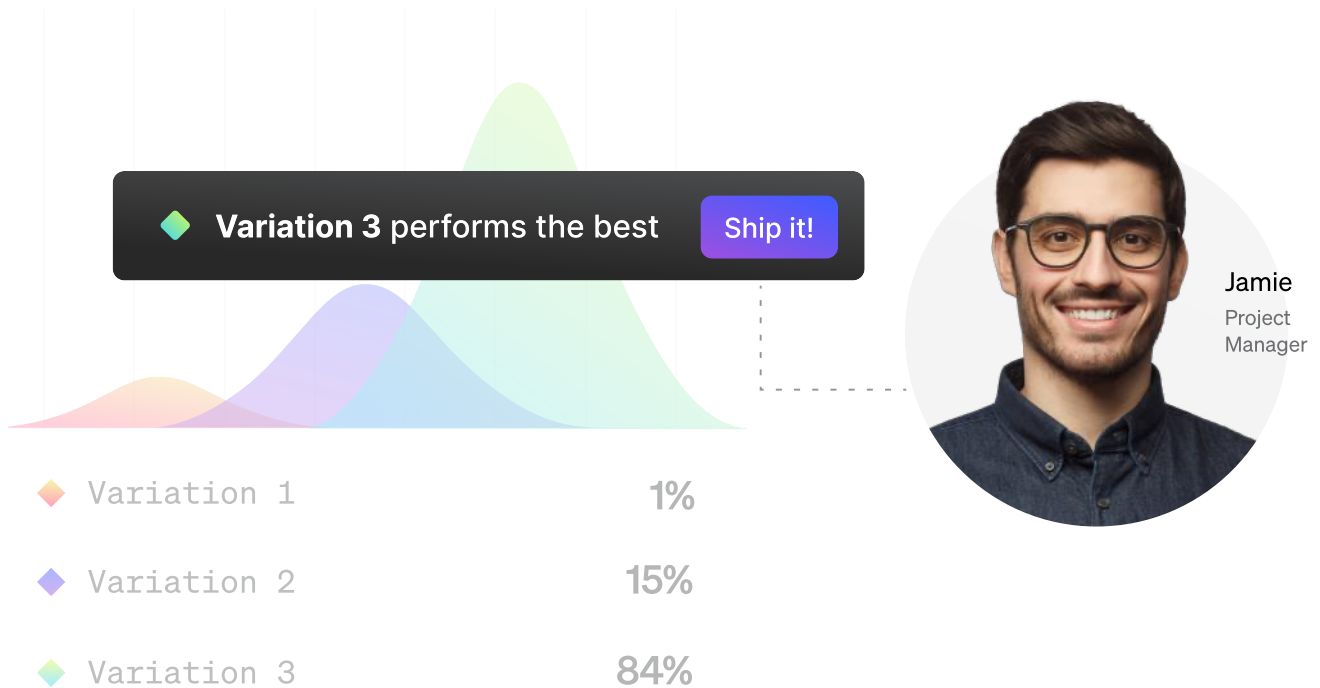
By gaining the power to continuously improve your products through experimentation, you can not only accelerate innovation but also ensure safety and scalability at every step along the way.

LaunchDarkly helps you uncover powerful insights and put them to work, illuminating every step on your path to strategic, scalable innovation. By effortlessly integrating with your preferred tools and workflows, LaunchDarkly makes it easier for every team to understand and act upon experiment results. You can export data to any tool of your choice, dissect it to your heart's content, and combine it with other data streams so your teams can see and interpret the big picture. For example, with LaunchDarkly's [Snowflake integration](#), you can export LaunchDarkly feature flag and experiment data to your Snowflake instances. And then you can analyze the data as you see fit.





No more crossing your fingers and hoping for the best. With LaunchDarkly, you can standardize data-driven product changes, ensure your decisions align with your goals, and foster a culture of continuous testing and learning that drives safe innovation at scale.

Finally, LaunchDarkly streamlines and democratizes experimentation, enabling stakeholders across teams to easily adjust experiments and take decisive actions. Once a winning variation of an experiment emerges, any team member can seamlessly roll it out to all users by toggling a switch—with no need to bother or wait for a developer.



Customer success stories



CCP Games unlocks the experimentation cheat code

Challenge

Makers of the wildly popular *EVE Online*, CCP Games struggled to connect its releases to real business impact as the company grew and releases became more frequent. The company was limited in their ability to run product experiments.

Solution

LaunchDarkly enabled CCP Games to run experiments that determined which feature variations were driving the most engagement. The solution allowed CCP Games to deliver, control, and run experiments on features in one seamless workflow. Plus, LaunchDarkly enabled those beyond the data science team to run, interpret, and act on experiments.



Results

A new and improved recommendation engine led to a rise in user engagement through longer median session length and more personalized experiences. The company's engagement with LaunchDarkly also informed its product roadmap for future enhancements, such as the creation of the AIR Career Program—an entirely new feature that helps increase user engagement.

//

LaunchDarkly has enabled self-serve experimentation at CCP. You don't have to be a data scientist to run valid, actionable experiments. This is unbelievably powerful.”

Nick Herring

Technical Director of Infrastructure, CCP Games



Read: [Full CCP Games case study](#)



Ritual

Ritual experiments its way to healthier results

Challenge

Before LaunchDarkly, Ritual used another feature management solution for feature flagging and running experiments. The solution did a poor job of integrating flags and experiments, creating a lot of extra development work.

Solution

With LaunchDarkly, Ritual can leverage feature flags and run experiments in one place. They can quickly launch experiments, clean them up, and release the winning feature variation of an experiment. They can enable a winning feature for the right audiences without having to make any code changes. And their teams have better visibility into the impact of new features.

//

...the time it took to not only implement a feature flag but also decide whether to experiment on that feature, re-implement in code how we would switch on the experiment, target audiences, and so on was considerable.

Daniel Archer

VP of Engineering, Ritual



Read: [Full Ritual case study](#)



Results

Previously, Ritual was only able to run 1-2 experiments per month. Now, they run 5+ per month. With feature flags and experiments in one place, the product delivery team is more agile, data-driven, and collaborative. Ritual uses flags and experiments for a host of critical use cases like conversion and retention optimization. In this way, LaunchDarkly positively impacts the bottom line.

//

I think that the main benefit of using LaunchDarkly experimentation is that it's directly tied to the feature flag implementation. It's not a separate implementation every time we want to test a feature.

Daniel Archer

VP of Engineering, Ritual



Get started

Ready to maximize the business value of every product feature? By replacing guesswork with data-driven insights, connecting metrics to features, and simplifying the process of deploying and running experiments, LaunchDarkly paves the way for more effective, efficient, and customer-centric product development strategies. Get started today by requesting a customized demo below.

[Get a demo](#)

