

LightStep [x]PM Enables Lyft's Move to Microservices, Helping Drive Significant Revenue and Improving Product Efficiency

Lyft is a ridesharing company, based in San Francisco and launched in 2012, that develops and operates its own mobile transportation app. Valued at \$11 billion and available to 95% of the population in the United States, Lyft is the fastest growing on-demand transportation service in the country.

The Challenge

Moving to Microservices to Support Exponential Growth

Lyft's consumer mobile app has real-time transactions totaling more than one million rides per day, so performance is critically important. The smallest lapses – even a few milliseconds – contribute to negative customer experiences and lost revenue. As Lyft's Vice President of Engineering, Pete Morelli, explained, "The bigger you get, the better you have to be. Half an hour of downtime may have cost you five rides early on, now it costs millions of dollars in rides. The level of reliability expected of Lyft is not trivial. People are riding to work or to doctors' appointments."

LightStep has played a critical part in helping Lyft minimize downtime and ensure that rider requests are quick and optimally routed.

Evolving Data Needs of a Microservices-based Architecture

In order to rapidly scale its system and support growth, Lyft started to explore moving from a monolithic architecture to microservices. Today, Lyft deploys more than 200 microservices in its distributed architecture, and this number is growing. These services work together to perform fundamental functions of the Lyft app, including matching riders with drivers, optimizing the route for the most efficient ride, and processing riders' payment information. It's a challenge to quickly and accurately monitor Lyft's system as the number of microservices grows, because a distributed architecture generates exponentially more data than its monolithic predecessor. This makes it more difficult to diagnose the root cause of a performance problem. To pinpoint the root cause of a performance issue, developers must understand every service that is invoked in the process of fulfilling a customer's ride request and the performance of each of those services. To gain insight into this detailed level of performance as efficiently as possible, Lyft chose to implement [x]PM as well as Envoy, the service mesh technology.



Challenges

- › Provide the fastest, best customer experience to each and every app user
- › Manage exponential growth and the increasing impact of long-tail latency on revenue
- › Observing 100% of application performance data across 200+ services
- › Connecting the dots from mobile to the backend to identify potential performance bottlenecks

Business Results

- › Improved efficiency of customer ride routes and accelerated response times by 60% (250 milliseconds)
- › Lowered Root Cause Analysis Time by 60%
- › Improved the postmortem process which resulted in fewer outages
- › Ensured end-to-end performance management across mobile, microservices, and monolith systems

Organization Details

- › Headquarters: San Francisco, CA
- › Industry Segment: Information Technology
- › Employees: 1,600
- › Funding: \$2B+

The Solution

Plug and Play with Envoy, the "Service Mesh"

Lyft was able to quickly and easily deploy [x]PM because it seamlessly integrated with Lyft's sidecar proxy or "service mesh", Envoy, which is a dedicated infrastructure layer that makes service-to-service communication safe, fast, and reliable. With the integration, every Lyft engineer automatically has access to [x]PM end-to-end distributed traces without any extra effort. Implementing Envoy and [x]PM has allowed Lyft to have state-of-the-art monitoring, and because [x]PM is a SaaS solution, it requires no maintenance from Lyft's engineers.

Diagnose Anything by Seeing Everything

The Lyft engineering team needed a solution that monitored 100% of its data. [x]PM is the only solution on the market that offers that capability. With its unique architecture, [x]PM computes advanced statistics based on all of the application performance data going through the system and stores examples of important information forever. As Morelli stated, "With [x]PM, there is no risk of overlooking any problems at the edges where the biggest problems are found. Look at it from a customer perspective – you might have a hundred rides that are awesome, but that one ride that was late or too slow, that's the one you'll remember. There's a bias toward that negative experience. When Lyft doesn't work, consumers go to taxis, competitors, buses, bikes, or to the car in the garage. They don't try again later."

"With [x]PM, there is no risk of overlooking any problems at the edges where the biggest problems are found. Look at it from a customer perspective – you might have a hundred rides that are awesome, but that one ride that was late or too slow, that's the one you'll remember. There's a bias toward that negative experience. When Lyft doesn't work, consumers go to taxis, to competitors, to buses, to bikes, to the car in the garage. They don't try again later."

Pete Morelli, Vice President of Engineering, Lyft

The Results

Saved Time + Saved Money = Happy Customers and Developers

[x]PM helps Lyft meet its goals of developer efficiency, system performance, and reliability by providing a complete picture of the software system. This allows engineers to investigate and resolve issues quickly, resulting in significant cost savings per year. Monitoring and application performance insights from [x]PM also empower engineers to make many critical-path optimizations that improve ride request times, increase dispatch efficiency, and ensure effective incident postmortems – all of which translates into increased revenue and developer efficiency.

According to engineers Roy Williams and Danial Afzal, one of the first projects where they used [x]PM was a spring cleaning of the entire system. The focus was on identifying and optimizing critical paths for dispatch services that connect riders with drivers. Saving time is a key goal, explained Williams: "The more time we get, the more efficient we can be. If we can use those extra milliseconds to find a more efficient match, that's a win for us, that's a win for our customers."



LightStep [x]PM Enables Lyft's Move to Microservices, Helping Save Significant Revenue and Improving Product Efficiency

Using [x]PM, engineers identified multiple unnecessary calls, removed them, and optimized the critical path for dispatch call services – saving significant cost and improving call time by up to 60% (250 milliseconds). Similarly, they found that Stripe, the payment authorization system, was set up to run serially. By parallelizing the process, they saved three seconds of precious customer time, improving performance by 60%. Afzal confirmed, “This is the first time that engineers have visibility, so they can take the necessary steps to ensure performance optimization and system efficiency.” In addition, [x]PM is critical for root cause analysis of performance degradations, debugging, and post-mortems of incidents. Afzal said, “Engineers are now able to efficiently understand a performance issue in a few minutes versus attempting to parse through logs, which can take up to an hour.” [x]PM is helpful for analyzing problems after they occur, but it is also useful to proactively prevent incidents. An outage incident can be devastating for a business, costing significant revenue. No business can afford it. [x]PM helps prevent such incidents from recurring by giving engineers the tools to analyze an incident and identify the root cause.

Lyft stays ahead of its continuing growth by providing engineers with visibility of the entire system, so they can resolve performance issues quickly. With [x]PM, engineers spend less time on debugging and incident responses, and have more time to build features that improve the usefulness of the app. That focus results in an enhanced end-user experience, more rides, and more revenue. As Pete Morelli says, “LightStep is the future of monitoring and was instrumental in our move to microservices.”

“LightStep is the future of monitoring
and was instrumental in our move to
microservices.”

Pete Morelli, Vice President of Engineering, Lyft

LightStep's mission is to deliver insights that put organizations back in control of their complex software applications. Its first product, LightStep [x]PM, is reinventing application performance management. It provides an accurate, detailed snapshot of the entire software system at any point in time, enabling organizations to identify bottlenecks and resolve incidents rapidly. LightStep is backed by Redpoint and Sequoia and is headquartered in San Francisco, CA. For more information, visit <https://lightstep.com> or follow at @LightStepHQ.

Learn how LightStep can help your organization measure and improve performance where it matters most.