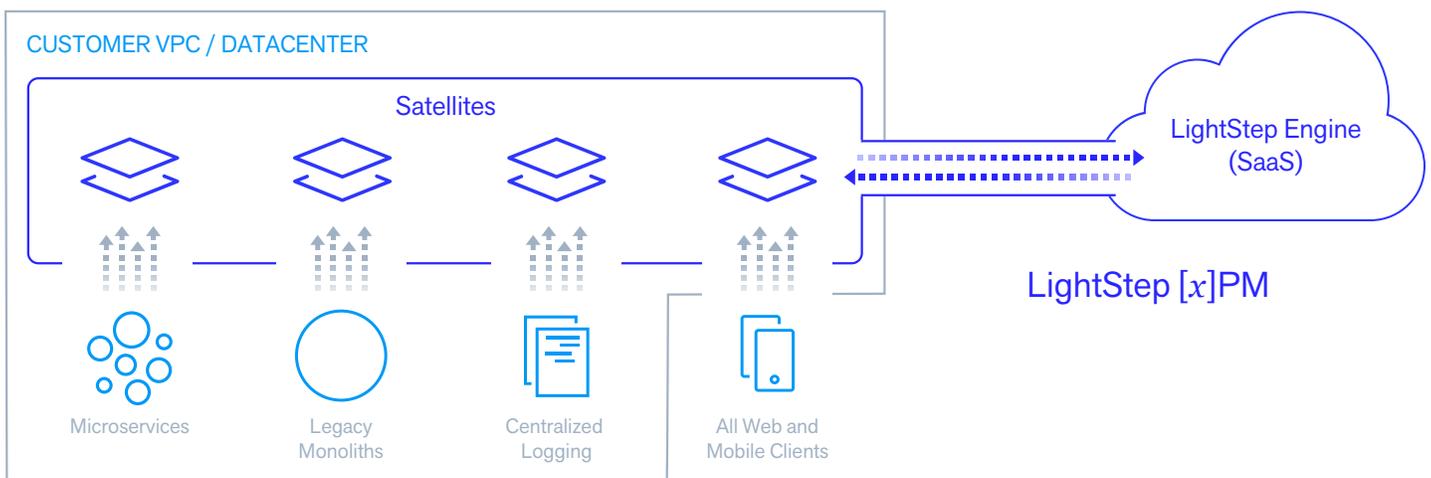


LightStep [x]PM

Engineering organizations are embracing the trend of designing new applications and decomposing existing, monolithic applications into modern microservice-based software systems that are built using various tools, frameworks, and cloud services. As these highly-dynamic applications become increasingly complex and distributed, it's vitally important for engineering teams to understand how the separate services and components interact with each other when diagnosing critical failures.

Satellite Architecture Overview

LightStep® [x]PM provides application performance management for the microservices era. It is the only platform that provides end-to-end tracing and metrics for complex transactions spanning every layer of an application, through monoliths and microservices, web and mobile clients. [x]PM monitors 100% of unsampled transaction data and is always-on in production environments, with negligible overhead. These features are possible because of [x]PM's unique platform architecture, which has two primary components: the Satellites and the LightStep Engine (SaaS).



Satellites enable [x]PM to capture a near-limitless amount of data. These software appliances collect unsampled performance and diagnostic data from every service, client, and server that make up a complete application. In order to handle this high volume of data, the Satellites buffer the data in memory just long enough to compute statistics and for the LightStep Engine to request the real-time data it needs to create end-to-end traces. Satellites are designed to scale horizontally without any storage requirements, and they can easily be deployed in any environment that supports Docker images, AMIs, or Debian packages, including Kubernetes.

At the heart of [x]PM is its SaaS component, the LightStep Engine, which coordinates the assembly of end-to-end traces by correlating all of the collected data distributed across the Satellites. The LightStep Engine identifies and weaves related trace data together into meaningful point-in-time stories about the application – even if the data was produced asynchronously or across distinct service boundaries. The LightStep Engine considers every operation and intelligently assembles traces automatically for interesting events like errors or latency spikes, as well as traces representative of normal operating behavior. Once assembled, these traces are stored indefinitely and can be reviewed at any time. By considering all of an application's transactional data, the LightStep Engine will reliably detect one-in-a-million anomalies, unlike any other technology, and show everything that happens both upstream and downstream from the event.

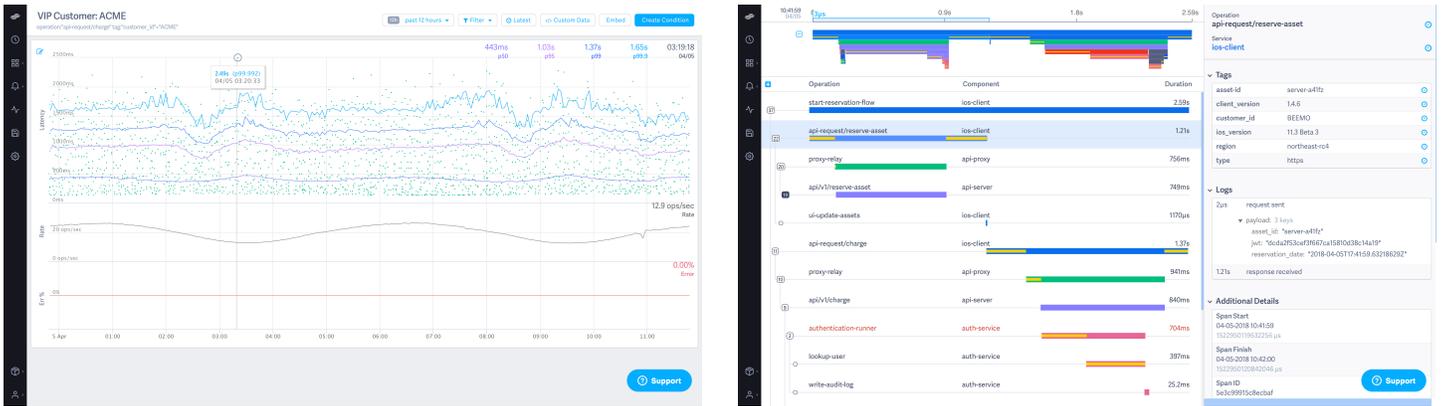
How Tracing Data is Gathered and Sent to [x]PM

Microservice-based applications generate huge amounts of data from all of the service interactions involved, so measuring and analyzing time-series and trace data for these systems is computationally difficult. Agent-based application performance management (APM) tools must do all of the analyses on the hosts themselves, so by design they are forced to sample data in order to limit this overhead. This sampling is particularly problematic for microservice-based applications, given the dynamic nature of the environment they operate in, and it often leads to major issues being missed entirely.

[x]PM gathers trace data from explicitly instrumented source code through open standards including OpenTracing, service mesh technologies (such as Envoy and Linkerd), open source adapter libraries, or even translations from structured logs that are centrally stored and managed. OpenTracing provides API libraries for distributed tracing whitebox instrumentation for many languages, including Go, Java, JavaScript, Python, Objective-C/Swift, Ruby, C++, and PHP. Because instrumented clients send data directly to [x]PM's Satellites without any analysis or preprocessing, this exchange is extremely lightweight and has virtually no performance impact on the hosts.

Intuitive and Feature Rich User Experience

[x]PM makes it easy for users to visualize, identify, and quickly resolve issues like anomalous latency spikes and system errors. [x]PM's web application interface allows users to drill into and examine the complex interactions for very large traces, across arbitrary time ranges and for any latency band, to uncover issues that may be buried deep in some backend service. It also automatically computes the critical path for each trace to allow users to jump straight to bottlenecks in the system, and it presents contextual log information and payloads inline for each operation.



In addition, [x]PM provides the flexibility to create Streams, or time-series trace data that match specific criteria defined by the user. The granularity of these Streams can be arbitrarily broad or focused without any cardinality limits, enabling enterprises to proactively monitor their system in the best way to manage their unique profile and priorities. For example, Streams can be used to record and compute time-series statistics and traces for all operations tied to a particular high-value customer or to compare the performance of related operations running through different mobile clients. Streams can be viewed concurrently in dashboards to make it easier to spot outliers and correlate patterns. Users can also filter high-volume, real-time data to immediately view traces for the aspects of their application that are holding back the system and the company's goals.

Lastly, [x]PM has native workflow integrations with Grafana, PagerDuty, and Slack, as well as support for JSON-based webhooks. These integrations make it easier for enterprise organizations to incorporate [x]PM's unique application performance insights into their existing workflows. Alerting through Slack and PagerDuty can be customized, so the proper team members are notified when arbitrary, user-defined thresholds are exceeded for certain aspects of the application. The evaluation window for these alerts is customizable from minutes to even weeks, so users can both detect immediate issues before they escalate and also monitor app performance for sustained customer trends that may have serious business consequences.

Enterprise-Ready Platform

- › Integrations with Cloud Native Computing Foundation (CNCF) technologies including OpenTracing, Envoy, gRPC, Istio, Linkerd, and Kubernetes.
- › Workflow integrations with Grafana, Slack, and PagerDuty.
- › OAuth support and Single Sign-On (SSO) with SAML.
- › Secure TLS communication between the LightStep Engine, Satellites, and Client SDK.
- › API to programmatically access, modify, and delete [x]PM's configuration and data.
- › Role-based access control (RBAC) to manage users.

Customer Successes

- › Twilio improves mean time to resolution (MTTR) by 92%.
- › DigitalOcean saves 1,000 developer-hours per month.
- › Yext reduces service level objective (SLO) violations by 70% for their top customers.
- › Lyft manages 100% of their application data, over 100 billion microservice calls daily, to reduce their customer ride request time by 60%.



Contact us by visiting lightstep.com/demo, and we'll show you how LightStep can help you gain control over your complex systems and dramatically reduce your MTTR.

Request a Demo

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.