

#### AN INFLUXDATA CASE STUDY

# Particle.io Users Iterate IoT Devices Faster with Data from InfluxDB





APRIL 2022

## Particle's Fleet Health Feature Provides Device-Specific Data For More Effective IoT Deployments

### Company in brief

Particle provides an integrated IoT Platform-as-a-Service that helps businesses connect, manage, and deploy software applications to connected devices, from edge to cloud and back. Over 240,000 developers, and 160+ Enterprise customers build on Particle, from fast-growing startups to Fortune 100 companies.

Particle's expertise goes beyond worldclass technology, enabling nextgeneration business intelligence, insights, and expert customer support to make sure IoT projects succeed.

#### Case overview

The Particle platform enables companies to connect, manage, and scale their connected products, bringing them to market quicker. To accomplish this, Particle developers needed to be able to collect telemetry data from a large number of edge IoT devices to measure performance. The company chose to use Telegraf and InfluxDB to collect and manage the time series data for the Fleet Health feature in its platform. With InfluxDB supplying granular, devicespecific metrics for Fleet Health, Particle's platform gives a broad user base the data they need to enable IoT capabilities on their products faster and more efficiently.

#### The business challenge

Companies often build IoT prototypes using microcontrollers like Arduino or Raspberry Pi. This means that when developers want to scale up, they have to rewrite their apps to work with field equipment. A key benefit of the Particle platform is its ability to scale from prototype to enterprise.

One significant challenge that Particle developers identified for their customers was accessing device-specific metrics. When rolling out large IoT deployments, having visibility into how each device performs is critical to success. Without this data, customers don't know what's going wrong with their devices so they end up spending more time with support, which slows down development. Giving customers direct access to their device data helps them identify issues quicker, or puts them in a better position when working with support to resolve issues faster. So, the Particle team needed a way to collect and present that granular data to their customers.



### The technical challenge



In the diagram above, the teal sections represent Kubernetes pods. Particle configured Telegraf as a sidecar for each pod to collect metrics through a Prometheus client for three different services: webhook service, API service, and device service. Telegraf then sends those metrics to InfluxDB. To get accurate information, users also need to know which devices are online. The Particle team created their Santa Claus service, which checks to see which devices are online. Another Telegraf instance collects that data and pumps it into InfluxDB, as well. Once the data is in InfluxDB, the Particle developers used InfluxQL, and later switched to Flux to run continuous queries on the data to restructure it to meet the needs of their platform, and to run scheduled tasks.

#### The solution



The platform shows the overall health of a device set.

The Particle team built the Fleet Health console in a mere six months. When selecting a group of devices, the console allows users to filter them by OS version, firmware, and time range, and presents metrics, such as the number of devices online, device connectivity and traffic, and success/failure data for function calls and other connection processes. Having access to this type of data allows users to quickly identify and isolate trouble areas. Particle also gives users API access to Fleet Health data.

# 66

Lying metrics are worse than no metrics. This is a really big deal. Generally, there are not going to be engineers reading these metrics; they're going to be product owners. They don't have the same tools at their disposal in order to understand that this graph is lying to them. You have to understand the pipeline in order to make sure that things aren't lying to you.

Cullen Murphy, Site Reliability Engineer, Particle

#### Results

While the goal of this project was to provide users with detailed information about their IoT devices, the Particle team realized that giving users too much data could also be a problem. Not everyone using the platform is an engineer or developer so they may not have the tools to understand what the data is telling them and whether that data is representative of what is actually happening.

Furthermore, enabling users to query anything at any time has the potential to slow down the continuous querying process. To combat this, the team found a way to provide data access in a way that serves everybody's needs without slowing down the system too much. To accomplish this, they developed ways to control bucket sizes in InfluxDB, and focused their efforts on a set of specific queries for that data. The result is a tool that's accessible for a broad range of users as they scale IoT deployments.

### About InfluxData

InfluxData is the creator of InfluxDB, the leading time series platform. We empower developers and organizations, such as Cisco, IBM, Lego, Siemens, and Tesla, to build transformative IoT, analytics and monitoring applications. Our technology is purpose-built to handle the massive volumes of time-stamped data produced by sensors, applications and computer infrastructure. Easy to start and scale, InfluxDB gives developers time to focus on the features and functionalities that give their apps a competitive edge. InfluxData is headquartered in San Francisco, with a workforce distributed throughout the U.S. and across Europe. For more information, visit <u>influxdata.com</u> and follow us <u>@InfluxDB</u>.

