AN INFLUXDATA CASE STUDY

Red Hat Relies on InfluxDB Platform to Create Single Source of Truth for Global Network Monitoring

NOVEMBER 2022

Red Hat Relies on InfluxDB Platform to Create Single Source of Truth for Global Network Monitoring

gNMI Data Collection Improves Monitoring Performance and Capabilities

Company in brief

Red Hat is a global leader for open source enterprise IT solutions. Its portfolio of products includes hybrid cloud infrastructure, middleware, cloud-native applications, and automation solutions. Red Hat is the largest open source company with 40,000 employees in 40 countries. More than 90% of Fortune 500 companies use Red Hat's products and services.

Technologies used:

InfluxDB, Flux, Kapacitor, Telegraf, Ansible, Grafana, PagerDuty, Splunk

Case overview

Red Hat's internal network monitoring team is responsible for monitoring the company's infrastructure. With offices all around the globe, the network monitoring team needs insight into 14,000+ interfaces and 1,600+ devices in 60+ locations to understand utilization and improve performance. Not all these devices use the same protocols. Some use older protocols, like SNMP, while others use more modern protocols like gNMI. InfluxDB gives Red Hat engineers broad compatibility to collect data from all these data sources and brings them together for a single source of truth. Using Flux language with InfluxDB also gives Red Hat engineers the ability to analyze and process their data in ways not possible before adoption, providing deeper insights and broader observability to the entire network. All this helps Red Hat's Ansible to automate more processes from its single source of truth, driving network efficiency and reliability.

> “Telegraf especially comes handy because it supports all of those protocols through gNMI, SNMP, but also the API. If that’s missing or there’s something that you need to add, you can create your own plugin, or you can just create a script and use the exec plugin. It’s very flexible. That was one of the reasons that we definitely wanted to go with Telegraf as the data collector for data acquisition.”

Martin Moucka, *Principal Network Engineer, Red Hat*

The business challenge

Managing enterprise IT infrastructure is a massive undertaking at Red Hat, which involves monitoring the backbone that supports over 40,000 employees in forty different countries. Red Hat’s internal networking monitoring team monitors over sixty of the company’s 105 global office locations. In total, that works out to be 14,000+ interfaces and 1,600+ devices.

Monitoring at Red Hat consists of performance metrics and visualizations. They want to know how their network infrastructure is performing, and visualize that data to better understand that performance. This also provides greater insight into all the pieces that comprise the network infrastructure and helps to detect blind spots and bottlenecks.

To achieve this type of observability, the network monitoring team sought to build a monitoring solution that functioned as a single source of truth for the network. To do so, they needed to be able to collect data from across the globe and make it available to the team in one location. This included data on device availability (e.g., ping, http, DNS), query speed, http response times and codes, external link utilization, latency, and more.

The technical challenge

With so many different interfaces and devices across the globe, the Red Hat team needed to be able to collect data from a wide range of data sources using the most efficient protocols available. They needed to visualize network performance, generate alerts, create network maps, and monitor network bandwidth, and needed a way to collect data to support these observability goals.

One of the biggest challenges is that the SNMP protocol remains very common in network monitoring. However, SNMP has several key limitations. One is that the protocol utilizes the device you're trying to collect data from when pulling metrics, which slows everything down. This meant that they couldn't pull metrics as frequently as desired and the smallest interval the Red Hat team could get SNMP metric data was one-minute. As a result, the team sought to switch to Google's Network Management Interface (gNMI) whenever possible. gNMI provides more granular data polling intervals, so Red Hat can get metrics every second from devices. gNMI has several other benefits for Red Hat, including to collect and store data types and metrics that SNMP cannot. It also enables Red Hat to write sensor data directly into the platform in some cases.

Not every device in Red Hat's environment supports gNMI, so how did the company bring everything together for its single source of truth?

The solution

Red Hat runs an enterprise instance of InfluxDB, which is a critical piece in their network monitoring architecture. They run Telegraf and Kapacitor on collectors distributed across the world. Red Hat uses Telegraf and the appropriate SNMP or gNMI plugins to collect data directly from network devices. They collect gNMI whenever possible but some devices only support SNMP, or are in the process of updating to gNMI support, so the data from these devices comes in via SNMP. For example, most interface utilization and error data come in as gNMI, but CPU and memory utilization are SNMP. BGP neighbors status is in the process of adding gNMI support so that data comes in as SNMP until that update completes.

Telegraf performs data enrichment when necessary before it goes to Kapacitor for analysis. Red Hat stores the analyzed SNMP and gNMI data in different measurements in InfluxDB and writes custom queries for each. InfluxDB allows them to combine this data at the query level, while keeping the data separate in the storage tier. If the Kapacitor analysis detects an issue, the system sends an alert to PagerDuty.

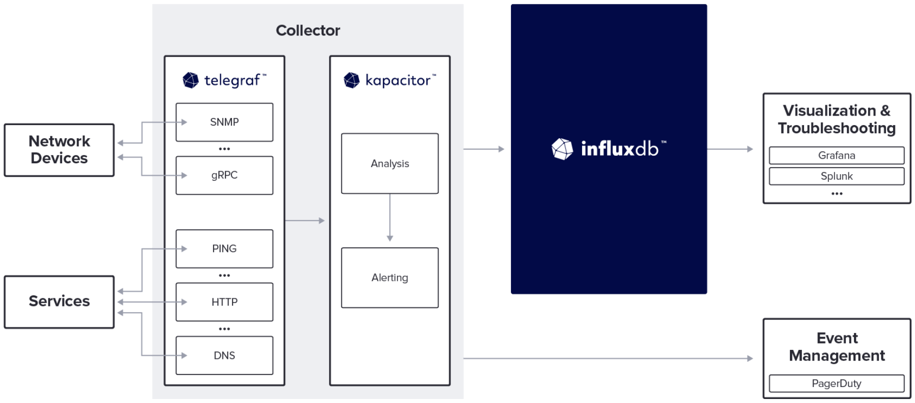

Red Hat uses the Flux language to query data in InfluxDB for visualization in Grafana and Splunk. Flux allows Red Hat engineers to connect different measurements and perform calculations between them in real-time. The ability to create functions in Flux and use them in queries also saves a lot of development time.

Dashboards generated from this data include a variety of information, like historic SLI/SLO data and real-time data visualizations.

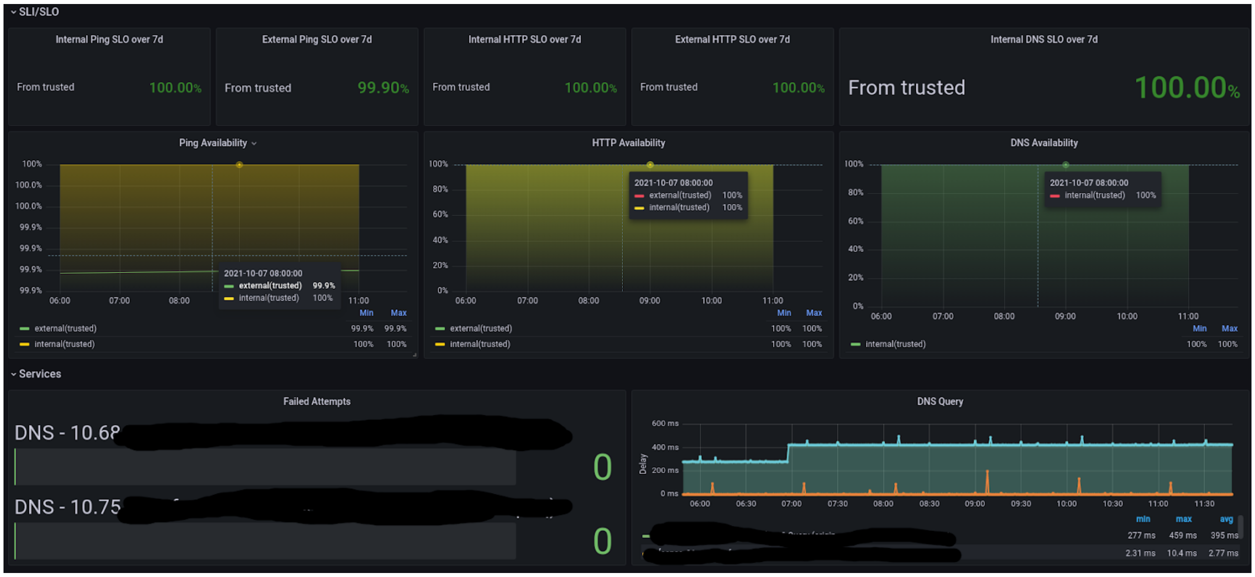

Results

The architecture diagram below shows the different components and data flows that comprise Red Hat's network monitoring solution. They rely on Ansible to handle network automation for device management and to configure Telegraf, Kapacitor, and InfluxDB instances.

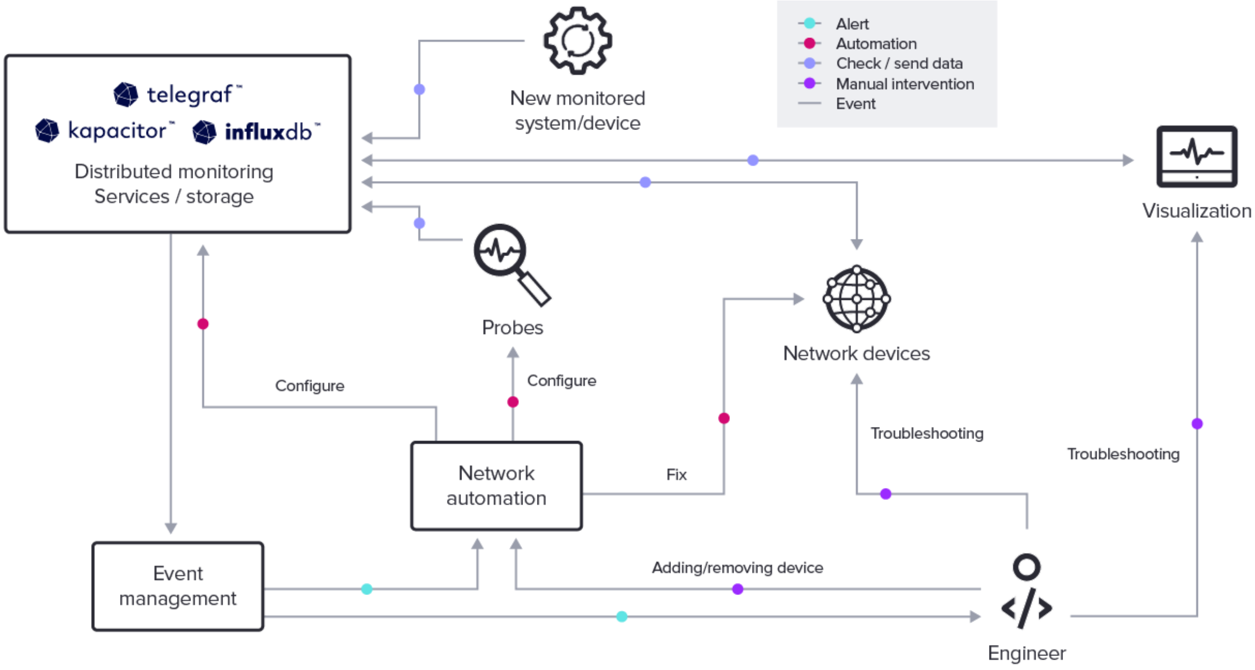

Thanks to this high degree of automation, this solution requires relatively little manual intervention, allowing those individuals to focus on critical issues, rather than managing individual devices and components. InfluxDB helps to provide those individuals with the broadest range of data possible to feed the single source of truth and facilitate automation, which improves real-time monitoring capabilities.

InfluxDB documentation, downloads & guides

Get InfluxDB

Try InfluxDB Cloud for Free

Get documentation

Additional tech papers

Join the InfluxDB community