influxdata®

AN INFLUXDATA CASE STUDY

PagerDuty Reduces Mean Time to Resolution with Process Automation and InfluxDB

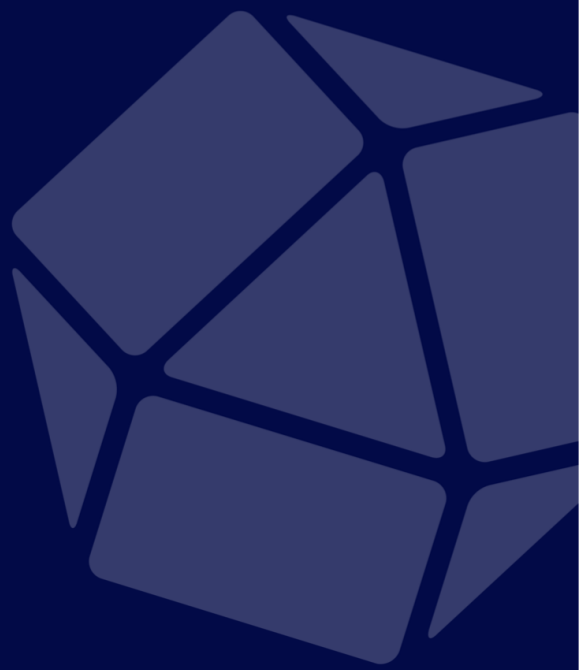

PagerDuty

NOVEMBER 2022

PagerDuty Reduces Mean Time to Resolution with Process Automation and InfluxDB

Time series data fuels performance monitoring and automation processes

Company in brief

PagerDuty is a digital operations management platform that helps businesses automate, orchestrate, and accelerate responses to critical incidents across their digital infrastructure. The PagerDuty platform alerts users about outages and disruptions to their IT systems. PagerDuty empowers users to take swift action for incident resolution to mitigate the effects of issues or to prevent them from occurring in the first place.

Technologies used:

InfluxDB, PagerDuty, Telegraf

Case overview

A key metric for service and support is mean time to resolution (MTTR). This is a measure of how quickly, on average, support teams identify, act, and resolve IT issues and incidents. The PagerDuty Process Automation team leverages InfluxDB to handle critical time series data that tracks infrastructure performance. This data triggers the PagerDuty platform's incidence management capabilities. PagerDuty, in turn, uses Process Automation to orchestrate, automate, and execute the different mitigation processes. This monitoring stack reduces MTTR and makes DevOps and SRE teams more efficient and effective.

> *"Whether I'm collecting from targeted databases, microservices, my network gear ... InfluxDB, has a plugin for that in Telegraf. And because it is designed around developers – it's a tool for developers – if there's something there that you want to tweak, update unique to your environment, it's very easy to do with the InfluxDB toolset."*

***Craig Hobbs**, Solution Consultant, PagerDuty Process Automation*

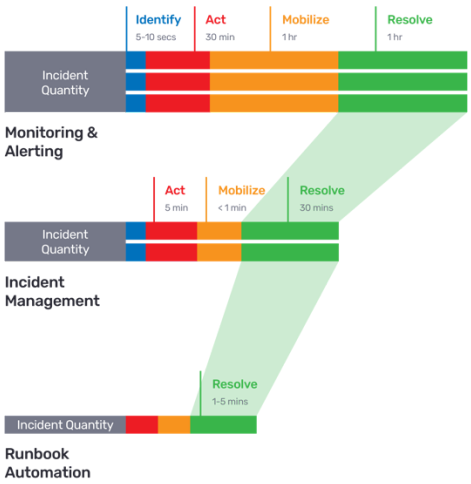

The technical challenge

When thinking about the incident management process, the PagerDuty Process Automation team breaks down the process into three distinct phases. Many different solutions aimed at reducing MTTR exist, but the PagerDuty Process Automation team sought best-of-breed solutions for each phase of the resolution process. At the monitoring stage, DevOps and SRE teams needed a solution that could handle the volume of data generated and, at the same time, cut through the noise of all that data to identify actual problems.

The next stage, incident management, required a tool to function as the nerve center of the resolution process. This needed to be able to aggregate the alerts, requests, and escalations coming in, to notify all relevant stakeholders, and to handle analytics and reporting on the resolution processes themselves.

Finally, the last stage, resolution, needed to be able to resolve issues and restore the production state. But, it also needed to handle post-mortem tasks, like updating incident management, which in turn updates relevant stakeholders so they can review processes and procedures. Any data generated during the resolution stage also gets rolled back into the monitoring solution to improve future incidents.

The solution

For each step along the way, the PagerDuty Process Automation team selected the best solution for that specific stage. They selected InfluxDB for monitoring, PagerDuty for

incident management, and PagerDuty Process Automation (formerly RunDeck) for resolution.

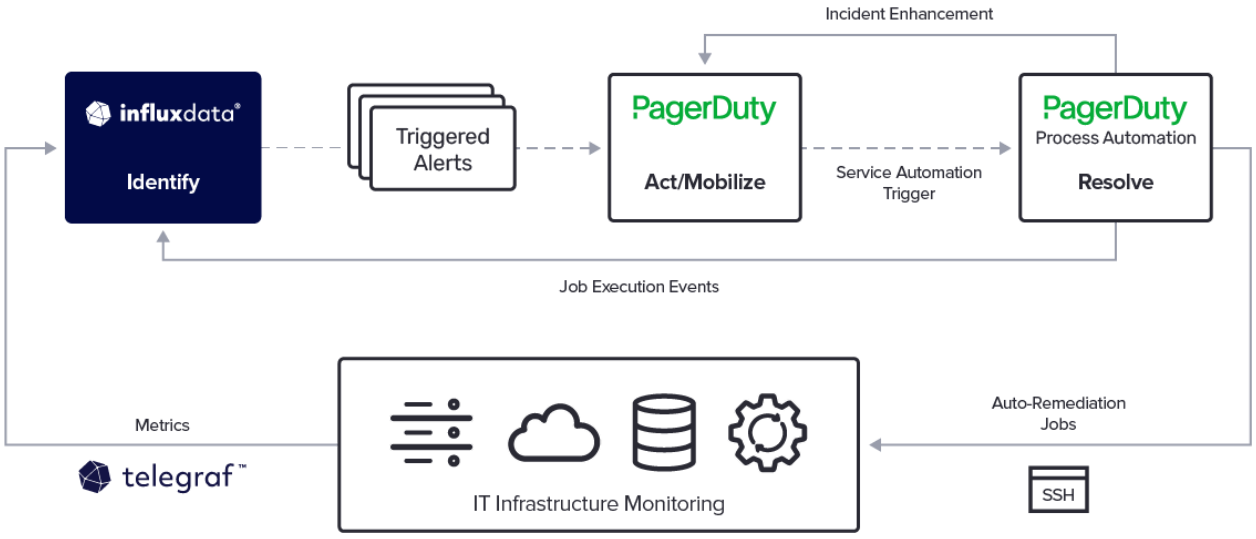

To collect infrastructure data, the team uses Telegraf. With hundreds of plugins, Telegraf can collect data from any infrastructure component, no matter where it's deployed or the environment it runs in. InfluxDB also supports templates, which can consist of modular tools designed to accomplish specific pieces of work. PagerDuty Process Automation team Solution Consultant Craig Hobbs built a custom InfluxDB template that allows the team to set up new monitoring instances, including Telegraf and all the necessary plugins, in minutes.

Once it's running, Telegraf sends all the data it collects to InfluxDB, the key technology in the monitoring stage. InfluxDB processes that high-volume time series data and intelligently dispatches all the triggers and alerts based on that data.

PagerDuty receives alerts from InfluxDB (and any other desired sources) and centralizes all the metrics and identifiers for issues. PagerDuty determines the appropriate runbook to resolve the issue and dispatches requests and services to PagerDuty Process Automation.

PagerDuty Process Automation serves as the runbook automation solution. It reliably executes the runbook assignments from PagerDuty. PagerDuty Process Automation maintains a complete activity trace history of every automation that it executes. This tells teams what ran, how long it took, what time it started, when it ended, and all the logging information associated with that job. PagerDuty Process Automation also returns data

about the executed runbook to both PagerDuty and InfluxDB, which can help both services refine the way they handle issues.

Results

These processes take place on the backend, so DevOps and SRE teams can use InfluxDB dashboards to track incident management in one place. Because PagerDuty feeds results back into InfluxDB, users can see when InfluxDB sends off alerts to the backend services as well as when the processes those alerts trigger complete.

The goal of this solution was to leverage best-fit tools that still provided developers with full control over what data gets pulled into the workflow and how to present that information visually.

Creating a consolidated view and workflow reduced the amount of time that DevOps and SRE teams spent working to resolve issues. Instead of escalating issues to personnel higher up the support chain, users have better visibility into resolution actions because they can quickly see system status and actively monitor resolution efforts.

By bringing together best-of-breed solutions, both open source and proprietary, PagerDuty created a solution that's flexible for developers and meets the demands of growing data volumes, various stakeholders and escalations, and the complexities involved in infrastructure monitoring and auto-remediation. It also helped reduce the number of incidents and overall MTTR. Utilizing InfluxDB templates makes the team more effective and efficient and is an innovative approach that other businesses can easily replicate because the modular nature of templates makes it easy to update or reconfigure them for new data sources or input tools.

InfluxDB documentation, downloads & guides

Get InfluxDB

Try InfluxDB Cloud for Free

Get documentation

Additional tech papers

Join the InfluxDB community