Scaling and Accelerating Your Kubernetes Journey with **Google Cloud and Splunk**

splunk> | Google Cloud

Very few open-source technologies have been adopted so widely — or in such a short time— as Kubernetes. Over the past five years, it's evolved from an internal project at Google to the world's de facto container orchestration standard. According to IDC, 70% of the world's enterprises will be using Kubernetes by next year.[1]

Today, companies are running Kubernetes environments at scale. This creates a demand for monitoring tools that can make sense of Kubernetes environments, to better support effective management, troubleshooting and data analytics.

We spoke with Splunk Product Marketing Director of ITOps, Amit Sharma, and Google Cloud Customer Engineer, Cuyler Dingwell, about the challenges DevOps teams face as they push the limits on scalability with Kubernetes. We also discuss the evolution of Google Kubernetes Engine (GKE), which builds on the company's legacy as the creator of Kubernetes and its experience deploying Kubernetes at scale. We also look at the recent release of Splunk Infrastructure Monitoring, which gives DevOps teams powerful new capabilities for monitoring and troubleshooting large-scale Kubernetes environments.

According to IDC, 70% of the world's enterprises will be using Kubernetes by next year.

1 IDC FutureScape: Worldwide Cloud 2020 Predictions, October 2019.

Q: Kubernetes was, of course, born as a project at Google. How has this legacy shaped Google Cloud's approach to scalability with Kubernetes, and how does your current Kubernetes solution address scalability?

A:

Cuyler Dingwell: One of the joys of being a customer engineer at Google is helping customers leverage cloud technologies like Kubernetes to handle massive growth. But I've also been able to take some lessons that are pretty unique to Google.

Today at Google, we spin up around four billion containers a week. That's over 6,600 containers every second! Leveraging those capabilities has allowed us to scale seven products to over a billion users, and Google Workspace has over two billion users. And that list doesn't include Google Cloud itself, which serves over a billion users.

In order to handle this type of scale, back in 2003, we developed Borg: a new type of cluster manager that replaced and unified a batch-oriented system and a service-oriented system. By 2014, Google had open-sourced Borg as Kubernetes and would also introduce it to Google Cloud as Google Kubernetes Engine. At this point, GKE has been in production use for more than five years, and in that time, we've seen our customers scale on GKE to handle massive workloads.

One of our most important focus areas for GKE involves helping customers develop secure, reliable and scalable applications without having to manage the underlying, supporting Kubernetes system. That's why GKE includes capabilities such as auto-repairing of Kubernetes clusters, automated upgrades and auto-scaling. In addition, we recently introduced a beta version of GKE auto-provisioning, which automates the process of creating Kubernetes node pools based on the type of workload a customer is running.

We've also added features to GKE to keep our customers secure as they scale their Kubernetes environments. We already strive to implement security best practices by default in GKE: We disable unnecessary services and enable encryption by default; employ a hardened host OS, container registry vulnerability scanning, and binary authorization to ensure the code running in your Kubernetes environment is known code; and support private Kubernetes clusters. But we also have new features, such as a GKE sandbox for workload isolation, and a workload identity feature that authenticates pods outside the Kubernetes cluster.

Q: Monitoring is obviously critical to understanding Kubernetes health and performance at any scale. What kinds of challenges are DevOps teams running into with monitoring as they deploy Kubernetes at scale?

A:

Amit Sharma: Clearly, challenges remain around complexity and monitoring in Kubernetes, especially at scale. And monitoring is more likely to become a top challenge as the size of a company grows — a bigger company needs bigger DevOps teams, and they implement a greater variety of Kubernetes application runtimes, languages and other components at scale.

Issues with complexity also contribute to performance and reliability issues with traditional monitoring tools. As a Kubernetes environment scales, for example, metrics with dimensions that have many different values, such as those using pod labels, can cause problems for monitoring tools that struggle with such high-cardinality metrics. As companies continue to scale their Kubernetes environments, these issues with monitoring performance and complexity can quickly get out of hand.

We must also keep in mind that Kubernetes is a distributed system with different components on the control plane and also on the worker nodes. When we deploy a service or application pod, many other things happen concurrently, and all of these underlying components must also churn and redeploy. As your Kubernetes infrastructure scales, so will the amount of churn in the environment.

Q: Are there other areas where conventional monitoring tools tend to struggle in a modern Kubernetes environment?

A:

Amit Sharma: There are several areas where traditional approaches to monitoring can fall short. First, scalability in Kubernetes is a multi-dimensional process, and traditional approaches to monitoring don't address this multi-dimensionality, resulting in inaccurate alerting or slow visualization.

Traditional monitoring tools also run into challenges related to data management and analysis in Kubernetes environments. Kubernetes monitors data flows from multiple sources and logs — all of them siloed, with each silo potentially using a different monitoring tool. Site Reliability Engineers (SREs) will have to constantly switch context as they move between tools to triage and troubleshoot problems, which leads directly to longer mean-time-to-resolution (MTTR).

Finally, with so much data coming in from so many sources, at such high velocity, correlating this data manually across multiple data sources just does not work. Traditional tools still don't leverage AI and ML capabilities to identify performance anomalies, and that's a huge problem.

Q: What should monitoring tools be doing to get ahead of these challenges and work effectively with Kubernetes at scale?

A:

Amit Sharma: We need a new approach to monitoring Kubernetes and Kubernetes-based applications. First, we need to understand the health of a Kubernetes cluster and all associated resources. Second, given the large number of components that get deployed or redeployed along with Kubernetes applications, we need visibility into the interdependencies between these components and their impact on the health and performance of the Kubernetes environment.

Granular logging is also important. The logs coming in from your containers, Kubernetes applications, metrics and monitoring platform, and other sources should include the context needed to support effective triaging and troubleshooting.

Next, your monitoring system needs to be able to handle high-volume, high-velocity data flows in real time. That means providing streaming analytics and accurate alerting no matter where the alert or the performance anomaly originates.

Finally, to do this, a Kubernetes monitoring solution must use AI-driven automation — bringing together machine learning, statistical modeling, and AI to automatically surface issues.

Q: How did Splunk approach the task of creating a Kubernetes monitoring tool that solves these challenges and supports scalability?

A:

Amit Sharma: We recently introduced Kubernetes Navigator (KN), a feature within Splunk Infrastructure Monitoring. KN gives DevOps and SRE teams a new way to manage and understand the performance of containerized applications. KN brings immediate value for teams that are just getting started on their cloud native journey, yet it also solves monitoring challenges for the world's most complex Kubernetes environments at scale.

KN uses an intuitive UI with hierarchical navigation that lets teams understand the performance of their entire Kubernetes environment. KN monitors the entire Kubernetes stack and uses dynamic cluster mapping to understand the relationships between dynamic components and detect interdependence issues.

KN also helps DevOps teams focus on the source of a problem with a granular, component-level view — moving from nodes to pods to containers and finally to workloads. Teams can pivot seamlessly to view applications, Kubernetes, and container logs, and correlate performance across the entire stack without losing context or switching tools. Then, KN uses AI-driven analytics to show teams the "why" behind performance anomalies, giving them the insights and recommendations they need to determine root causes.

Finally, Splunk is a unified platform that ensures a single source of truth for understanding Kubernetes performance. A highly scalable, streaming metrics architecture, with separate databases for human-readable metadata and time-series metrics, supports instant visibility. Charts and dashboards in Splunk refresh every second, and alerts and automation tasks trigger within seconds of an event. Existing solutions in the market take as long as five or 10 minutes to perform similar tasks.

Kubernetes scalability:
An increasingly urgent issue for DevOps teams

Kubernetes scalability is likely to become an even more urgent issue for DevOps teams as companies continue to deal with growing demand for digital offerings, and as teams look for better ways to innovate, accelerate, and deliver unique customer experiences. This trend will make challenges using Kubernetes at scale increasingly costly and disruptive, and make it harder for these companies to stay competitive and to deliver great customer experiences. And it's a trend that will make solutions like Google Cloud GKE and Splunk Infrastructure Monitoring increasingly valuable for teams that need a practical and predictable path to implementing Kubernetes successfully at scale.

Learn More.

Understand and manage the performance of your Kubernetes environments now — faster than ever before — with Splunk and Google Cloud.

Get Started

21-17885-Splunk-Scaling and Accelerating your Kubernetes Journey with GCloud and SPLK-EB-107