

# Solvvy

## EMPOWERED BY SCALYR™

### CASE STUDY

## ABOUT

# Solvvy

Solvvy is reinventing how customers interact with businesses via its innovative and intelligent conversational platform. Powered by artificial intelligence, the Solvvy conversational platform scales support operations with intelligent self-service so businesses can improve customer satisfaction and meet ever-increasing customer expectations while simultaneously decreasing costs.

## OBSERVABILITY CHALLENGES

As part of their infrastructure modernization efforts, Solvvy had recently migrated from Amazon Web Services (AWS) to the Google Cloud Platform (GCP) and moved entirely to Kubernetes. On AWS, the team had limited visibility into their system, so they wanted to take this opportunity to improve that.

Solvvy needed an observability solution that could effectively handle all of their log data and enable their development and operations teams to easily and effortlessly search through it when debugging problems. Although StackDriver was available to them, they knew that they needed a modern solution that offered a more robust set of features, particularly with Kubernetes. Solvvy initially tried to leverage an existing in-house ELK solution but found that its Kubernetes support was limited and it would be challenging to learn, configure and maintain at scale. This led to a decision to investigate and evaluate managed solutions that would scale gracefully with their conversational platform and, most importantly, be easy to set up and maintain in their Kubernetes environment.

## WHY SOLVVY CHOSE SCALYR

Solvvy approached their evaluation of managed observability solutions with three main criteria:

- 1 Minimal effort to set up and maintain,
- 2 Robust Kubernetes integration and
- 3 Affordability

Sumo Logic was disqualified immediately on price. While their initial shortlist included managed Elasticsearch, StackDriver, Datadog and Scalyr, their evaluation criteria quickly whittled the list down to two serious contenders: Datadog and Scalyr. Managed Elasticsearch fell short on Kubernetes integration. While StackDriver had good Kubernetes integration, dashboard build-outs and analytics were slow and inflexible.

Although Scalyr had an edge over Datadog with respect to Kubernetes support and price-to-value, driven by simple, transparent and inclusive pricing, the deciding factor in choosing Scalyr came down to a fundamental architectural difference that manifests in data ingestion and handling, and significantly impacts set up and maintenance. Unlike Datadog, Scalyr does not require indexing data, so users have immediate access to their data and can easily search and filter on any attribute defined by configured parsers or appear in json logs. Datadog uses keyword indexing, which forces users to explicitly add attributes to each index in order to do any meaningful analytics or filter on them, which is a non-trivial task for users and creates a layer of administrative and maintenance overhead that quickly becomes unscalable. The flexibility that Scalyr enables by not explicitly indexing let Solvvy get started immediately and helps them avoid any ongoing maintenance as the shape and structure of their log data changes over time. Scalyr's architecture also avoids the limitations that Datadog places on the maximum number of attributes.

## CUSTOMER

# Solvvy

## CHALLENGE

Needed a cost-effective and no-maintenance observability solution that effectively handles Kubernetes and application logs.

## RESULTS OF USING SCALYR

- ✓ Easy to set up - able to ingest and search all log data via Scalyr parsers
- ↔ No maintenance - developers have the flexibility and freedom to change code without impacting observability
- ➔ Increased troubleshooting and debugging productivity in Kubernetes environment

Scalyr has empowered Solvvy to execute and scale their business quickly to meet the needs of their growing customer base. Instead of spending time configuring and loading data into their observability tool, Solvvy engineers can focus on shipping code fast. Using Scalyr, Solvvy engineers have saved countless hours by avoiding the time-consuming and error-prone process of manually adding attributes to indexes, and have been able to identify and remediate root causes to complex problems in their Kubernetes environments within minutes. Developers also don't need to spend time worrying about the structure of their logs since parsers are easily configured within Scalyr, giving them the freedom to experiment, quickly create new services and delivering features at the fastest pace possible.

**“Being able to change code without needing to manually update indexes or the configuration of our observability tool is a beautiful thing.”**

- Christian Theilemann  
Software Engineer & Chapter Lead



**ABOUT**  
**SCALYR™**

Scalyr is a log management and observability platform for the new stack. Purpose-built to handle the scale and complexity of modern cloud architectures, Scalyr changes the dynamics of delivering healthy applications by allowing engineers to quickly troubleshoot problems and focus on doing what they love -coding. With 90% of searches completing in under one second and thousands of active users, Scalyr has transformed logs from afterthought into advantage. The company has the highest rating in its category in G2 Crowd, which is a Gartner 2018 Cool Vendor, and was recognized as a 2018 Forbes Cloud 100 Rising Star.