



Customer

trivago
www.trivago.com



Industry

Global travel search engine,
helping users find and compare
hotel deals.



Scale

Operating in 190+ countries
Handling millions of queries daily



Challenges

- Migrating from Google Cloud Run to Kubernetes without losing scale-to-zero functionality.
- Maintaining a cost-efficient preview environment for 150-200 concurrent pull requests.
- Ensuring unified deployment pipelines for logging, monitoring, and observability.



Overview

Kedify's solution provided trivago with a seamless migration path from Google Cloud Run to Kubernetes while retaining scale-to-zero efficiency, reducing complexity, and standardizing deployments.

trivago

How Trivago Migrated From Google Cloud Run to Kubernetes with Kedify

The Challenge

trivago had been running its infrastructure across both Kubernetes and Google Cloud Run, using Kubernetes for core workloads while relying on Google Cloud Run for test preview environments and certain other workloads. Google Cloud Run's scale-to-zero capabilities made it an attractive option for managing temporary environments, allowing multiple test instances to exist without incurring unnecessary costs. However, this split between platforms introduced operational complexities.

“

Google Cloud Run provided us with a very cost-efficient way to manage preview environments. Our main challenge was ensuring that Kubernetes could offer the same efficiency while integrating better with our standardized tooling.

— Jakub Sacha
SRE, Trivago

”

While Google Cloud Run provided a flexible and low-maintenance scaling model, having workloads spread between Google Cloud Run and Kubernetes led to fragmentation in deployment workflows, release processes, and debugging practices. Engineers had to manage different operational approaches depending on where a service was running, resulting in inconsistencies and additional maintenance overhead. Observability was also a challenge, as monitoring and logging solutions differed between both platforms, complicating troubleshooting and increasing operational complexity.

The company aimed to standardize its infrastructure on Kubernetes to improve overall operational efficiency. The key challenge was maintaining the scale-to-zero functionality for preview environments, as trivago regularly has **150-200 pull requests open**, each requiring its own test deployment.


Keeping these environments constantly active in Kubernetes would result in high resource consumption and operational inefficiencies. On average, these preview environments were only actively used for about one hour per day, meaning they remained idle for the vast majority of the time. Additionally, after a pull request was merged, the corresponding preview environment would remain on the cluster for up to two weeks to allow QA teams to revisit and verify issues against production—further increasing the need for an efficient scaling solution.

The Solution

trivago partnered with Kedify to implement an event-driven scaling solution based on KEDA and Kedify's HTTP scaler, which enabled them to keep their desired scale-to-zero functionality within Kubernetes.

Kedify's HTTP-based autoscaler allows test environments to be provisioned only when accessed. Developers could simply click a link in a pull request, triggering Kedify to dynamically spin up the necessary deployment for testing. Once the environment was no longer needed, Kedify automatically scaled it back to zero, eliminating unnecessary resource consumption. This ensured that trivago's infrastructure remained cost-efficient while maintaining the operational flexibility of Google Cloud Run.

Beyond just scaling, Kedify's solution played a key role in streamlining trivago's operational processes.



“

With Kedify, our developers get the best of both worlds, cost-efficient scaling like Google Cloud Run, but fully integrated within our Kubernetes-based platform.

”

— Jakub Sacha
SRE, trivago

With their entire stack now running on Kubernetes, they could apply a unified release process across all services, simplifying deployments and reducing overhead. Engineers no longer had to switch between different deployment models, which improved consistency and reduced friction in development workflows. Additionally, observability and debugging became more efficient, as every service could now integrate seamlessly with their preferred tools like Prometheus and Elasticsearch.

The Impact

trivago successfully transitioned from Google Cloud Run to Kubernetes while **maintaining the cost-efficiency of scale-to-zero functionality**. With Kedify's autoscaling, test environments are provisioned dynamically, ensuring resources are only used when needed. This optimization not only reduced operational complexity but also minimized unnecessary infrastructure costs.

Standardizing on Kubernetes enabled trivago to streamline resource management across the organization, enhancing scalability and maintainability. By eliminating the need for custom workarounds, developers can now focus more on innovation rather than infrastructure management, leading to a more efficient and agile engineering workflow.

Results



Scale to Zero Efficiency

Kedify enables trivago to scale deployments only when needed, reducing waste.



Unified Deployment Pipeline

Seamless integration with Kubernetes and standardized deployment tooling.



Seamless Migration from Google Cloud Run

Maintained flexibility and operational efficiency.

Please reach out for more information, to try a demo, or to learn more:
www.kedify.io.