AKAMAS

SUCCESS STORY

Sabre achieves optimal cloud performance and cost efficiency with Akamas

How Akamas accelerated Sabre's cloud migration and application modernization journey

Sabre®

Sabre

The technology that powers travel

Sabre Corporation is a software and technology company that takes on the biggest opportunities and solves the most complex challenges in travel. The Company **connects travel suppliers and buyers around the globe** and across the ecosystem through innovative products and next-generation technology solutions.

Sabre harnesses speed, scale and insights to build tomorrow's technology today – empowering airlines, hoteliers, agencies and other partners to retail, distribute and fulfill travel worldwide. Headquartered in Southlake, Texas, US, Sabre serves customers in **more than 160 countries** around the world.

Leveraging cutting-edge technology and a commitment to innovation, Sabre enables its customers to deliver seamless travel experiences to millions of travelers every day.

From airline reservation systems to hotel distribution platforms, Sabre's comprehensive suite of solutions empowers its clients to **optimize operations**, help drive revenue, and enhance customer satisfaction.

With a focus on sustainability and digital transformation, Sabre continues to drive the evolution of the travel industry, positioning itself as a trusted partner for businesses looking to thrive in a rapidly changing landscape.

50K+
travel agencies

11B+
shopping requests per month

200
countries & territories

$2.9B
total revenues*

*as reported for fiscal year 2023

© Akamas S.p.a., 2024. All rights reserved

AKAMAS

The Challenge of cloud-native app modernization

In the age of digital-driven business, incumbents that want to maintain their leadership need to react to market needs as quickly as possible. For this reason, Sabre embarked on a **cloud journey** and decided to modernize existing monolithic applications to microservices. This allows for fast and flexible software delivery which in turn provides value to the business and its customer quicker.

One of the most critical challenges that Sabre encountered in moving to a cloud-native environment, was the need to complete the migration quickly and without impact to customer experience. This in turn created significant cost efficiency challenges. Once those primary objectives were achieved, Sabre focused on reducing costs as much as possible by focusing on right-sizing applications.

Challenges emerged also with applications rearchitected into microservices. Transitioning from monoliths to microservices in the cloud demands a shift in mindset, particularly regarding cost efficiency and application performance. While Kubernetes provides an excellent platform for precise resource allocation and auto-scaling, proper configuration is essential.

Should you opt for smaller pods for quick and easy scaling, or medium-sized pods requiring less autoscaling? Similarly, when tuning JVM, do you use larger JVMs for less frequent but longer garbage collections, or smaller JVMs with more frequent yet quicker garbage collections?

An incorrect setup means failing to harness the efficiency advantages of the cloud and compromising application reliability SLOs.

"We decided to look for a solution to automate the optimization process in a repeatable way. Akamas is very intuitive and has the ability to run the optimization process very fast."

Pawel Popiolek , Senior Principal Software Engineer at Sabre

How Akamas helped

Sabre's decision to choose Akamas stemmed from its strategic imperative to reduce cloud costs. During the proof-of-concept phase, Akamas generated promising gains, improving resource utilization by 20% of a critical gateway application before the migration to the cloud.

Following their migration to the cloud, Sabre sought to expedite the process of right-sizing resources and reduce costs as much as possible. Several applications were onboarded to Akamas to facilitate quick and effortless identification of the most efficient configurations, tailored to various tech stacks.

For Java-based microservice applications, Akamas allowed Sabre to identify the best configuration for the JVM, including memory sizing and garbage collector types, and Kubernetes pod CPU/memory resources.

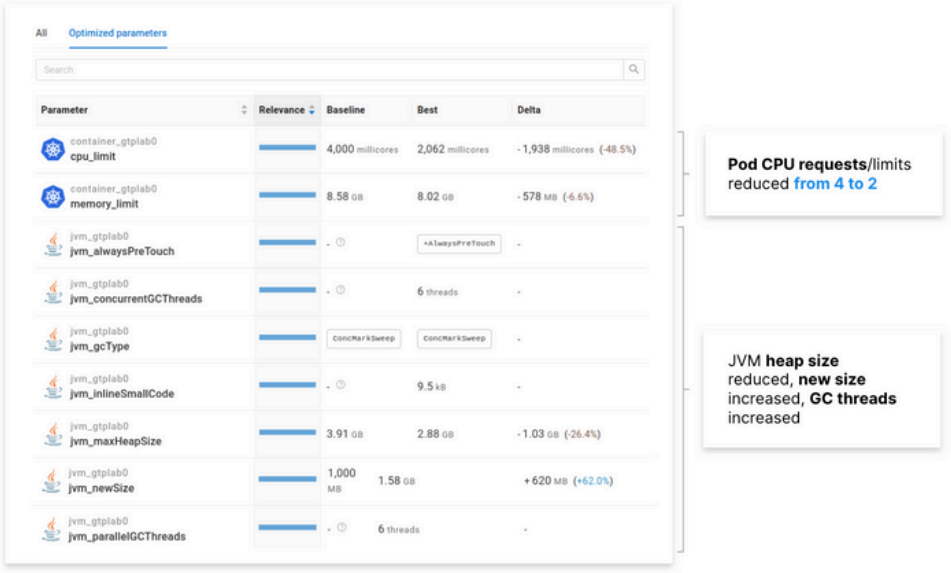

© Akamas S.p.a., 2024. All rights reserved

AKAMAS

Akamas drives the optimization based on goals and constraints, according to specified goals and constraints. The team's objective was to minimize costs while maintaining application response time and error rate **Service Level Objectives (SLOs)**, while also avoiding out-of-memory problems.

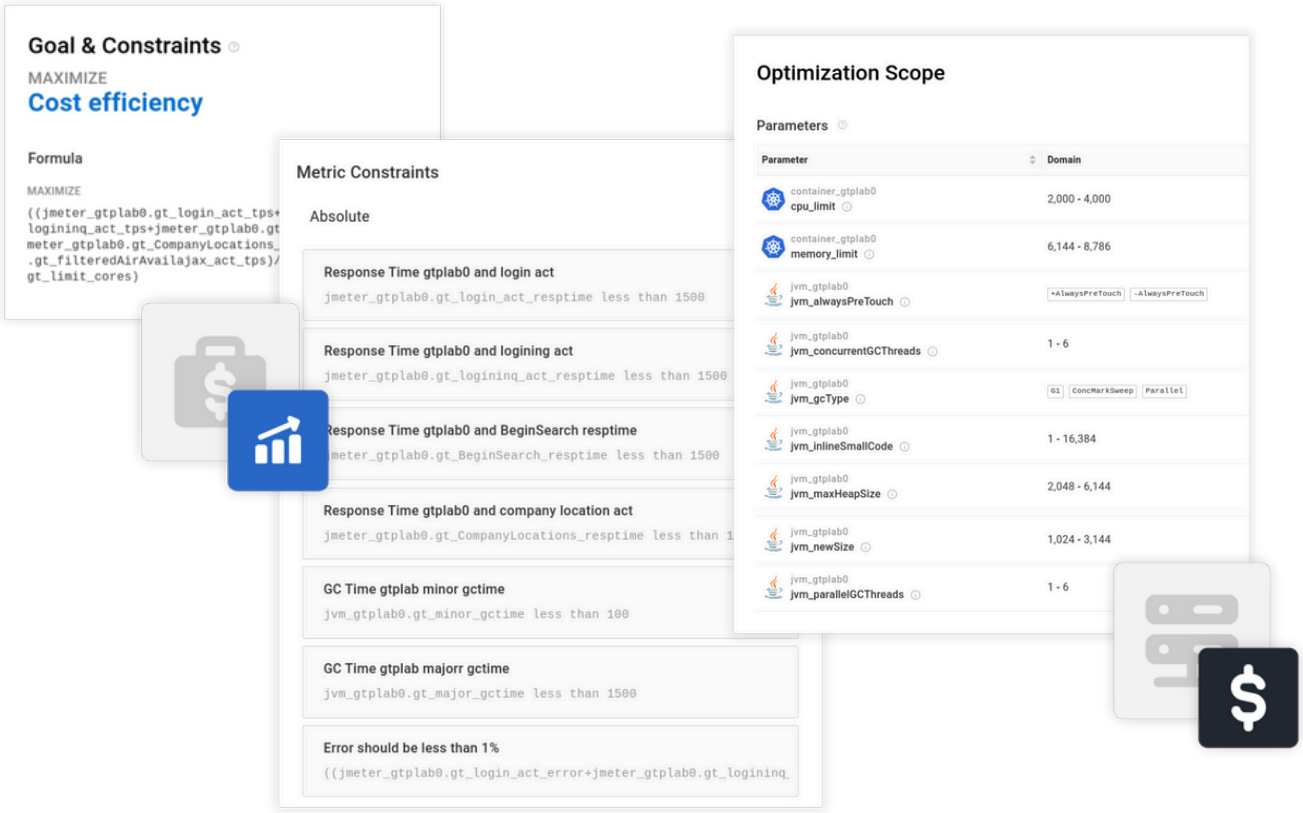

For monolithic applications, Akamas allowed the team to automatically identify the best GCP instances among the many families, series, and types available.

A key question was whether instances with fewer CPUs, or more affordable processor architectures, could reduce costs while still meeting Sabre's application performance SLAs. Akamas provided the answer, enabling the team to select the most cost-effective instance that still met the requirements for application throughput and response time.

"Whether you are right-sizing cloud infrastructure or tuning JVM performance, Akamas can save you tons of time and work."

Aodan O'Sullivan, Director of Performance Engineering at Sabre

© Akamas S.p.a., 2024. All rights reserved

AKAMAS

Benefits

For Sabre, Akamas was instrumental in quickly addressing many critical questions surrounding the optimal configuration of application and infrastructure stacks, empowering them to achieve optimal performance and cost efficiency for their cloud-native and monolithic applications. Sabre reaped numerous benefits from their collaboration with Akamas, both in terms of enhanced application performance and improved cost efficiency across their cloud-native applications.

Sabre achieved an approximate **50% cost reduction** for one application service recently migrated to Kubernetes.

Speed of optimization was another key benefit: the time required to optimize an application went **from weeks to a few days**. Akamas was in fact instrumental in the cloud migration and re-architecting of applications from monolith to microservices, saving Sabre significant time and effort in the process. Application development team often overlook, or do not have the resources for application cost efficiency and performance tuning. Akamas made this easy, simple, and quick.

Akamas also helped Sabre's skilled performance engineering team become better at their job. The team learned useful lessons and best practices that made everyone better. The team noticed some **patterns for improvements** for applications that have a similar profile that they were able to test and verify quickly. With Akamas, Sabre performance engineering team can better serve their internal customers: application teams can ship applications that are at the same time more cost-efficient and offer better performance to their users.

This strategic partnership with Akamas has the ability to enable Sabre to realize tangible benefits in terms of performance optimization, cost reduction, and overall operational efficiency.

50%

application
cost reduction

4-6x

optimization time
acceleration

Better

team collaboration
& best practices

© Akamas S.p.a., 2024. All rights reserved

AKAMAS

AUTONOMOUS OPTIMIZATION
FOR MODERN APPLICATIONS

AKAMAS

The AI-powered platform for live and offline application optimization

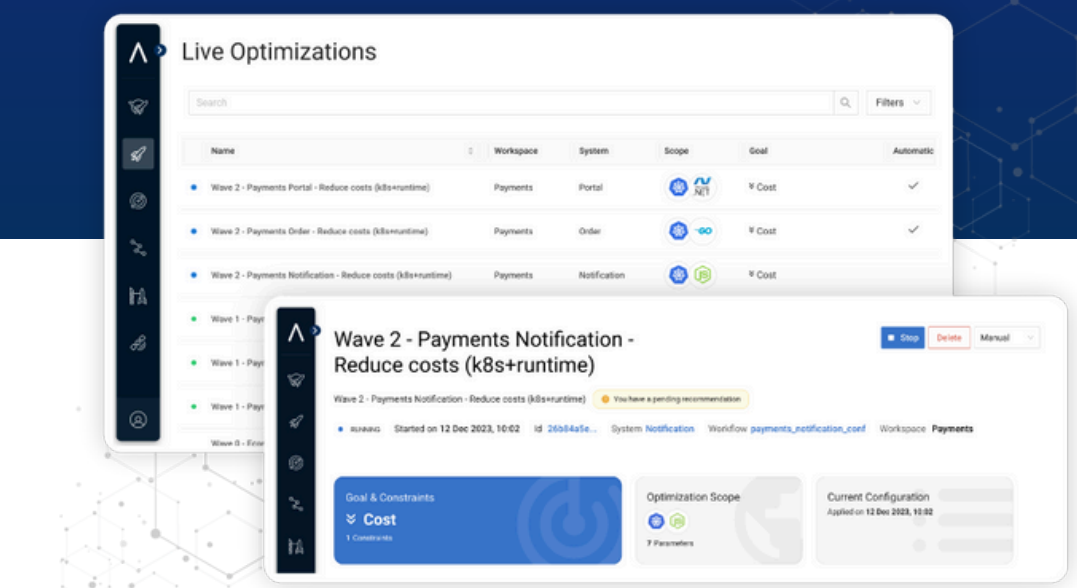

REDUCE
CLOUD COSTS
-60%

Cut your applications' demand for compute and infrastructure resources.

IMPROVE
EXPERIENCE QUALITY
+30%

Increase throughput and reduce response time, with lower fluctuations and peaks.

ENSURE
SERVICE RESILIENCE
zero

Ensure that apps and microservices work smoothly through workload peaks and anomalies.

INCREASE
TEAM PRODUCTIVITY
5x

Automate application tuning, cutting entirely time spent on manual configuration.

Sabre · lastminute.com · Sisal · RAIFFEISEN · sogei · TeamSystem · NAVAN

Milan HQ
Via Schiaffino, 11
20158 MILANO
T: +39 02 4951 7001

Boston
211 Congress Street
Boston, MA 02110
T: +1 617 936 0212

Los Angeles
12130 Millennium Drive
Los Angeles, CA 90094
T: +1 323 524 0524

Singapore
5 Temasek Blvd,
Singapore 03898

© Akamas S.p.a., 2024. All rights reserved

AKAMAS