AKAMAS

Autonomous Optimization for Kubernetes applications

Kubernetes is critical for app efficiency

Efficiency is one of **Kubernetes** top benefits, yet many companies often experience higher than expected **infrastructure/cloud costs** as well as **performance and stability issues**.

The **complexity of Kubernetes resource management** often leads developers, Performance Engineers and SREs to adopt very conservative configurations and **resource overprovisioning**.

The resulting **unnecessary infrastructure/could costs** may significantly affect the overall **cost efficiency** of delivered services, while not necessarily removing the **risks of missing SLOs**.

55%

of companies mentioning lack of in-house skills and limited manpower as the #1 challenge with Kubernetes

Source: Statista

46%

of Kubernetes users with the unmet challenge of ensuring application performance and resilience

Source: Akamas survey

68%

customers facing increasing Kubernetes costs and missing/inadequate capabilities to avoid overspending

Source: 2021 FinOps report

AKAMAS

Tuning Kubernetes is very challenging

Kubernetes cluster size is determined by **resource requests** set by developers in the container manifest for each application - so **clusters can be full even if utilization is very low** (e.g. 1%).

Sizing Kubernetes containers/pods by properly **setting resource requests and resource limits** may require weeks or months **for each single microservice** - so **overprovisioning** is common.

Kubernetes autoscaling enable automatic scaling under load but can also **magnify latency & efficiency issues** when Kubernetes microservices applications are not well configured.

too low resource requests
may cause slowdowns and reliability issues - **too high requests** may lead to wasted resources and overspending

too low resource limits
impact application stability and performance as containers get killed or slowed down under resource pressure

too high resource limits
allow runaway containers to act as “noisy neighbors” and degrade performance of other applications sharing the same nodes

AKAMAS

Autonomous Performance Optimization

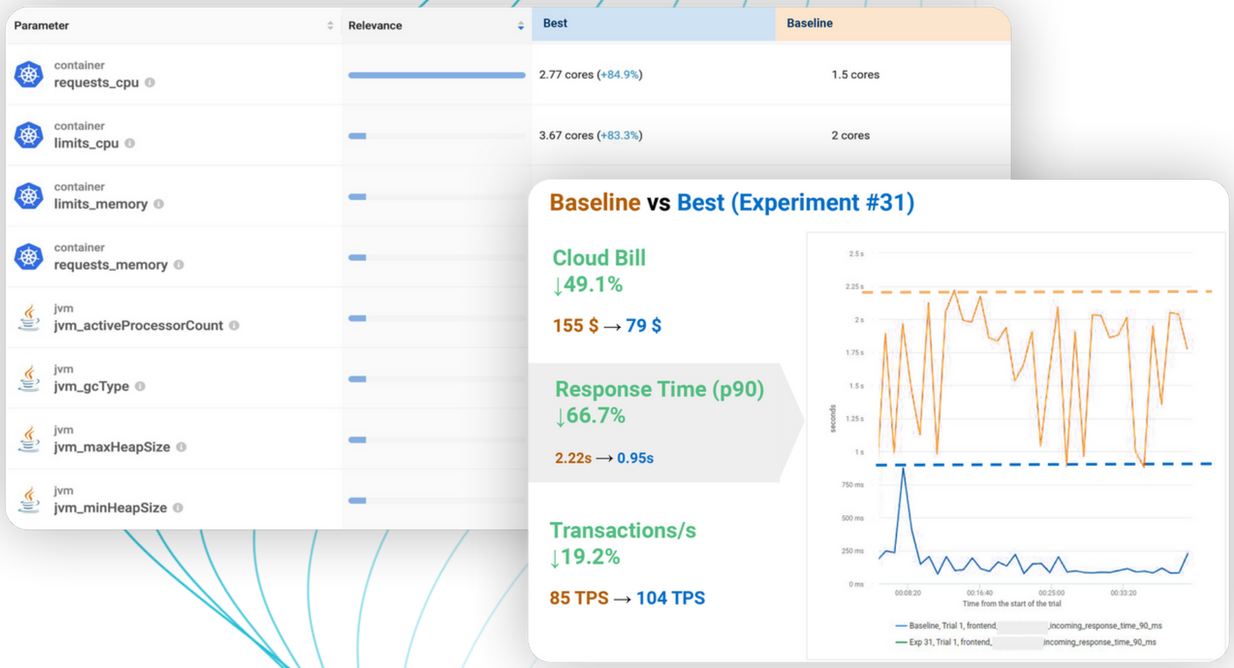

Akamas AI-powered optimization **automatically identifies the best configurations** for hundreds of parameters for **any Kubernetes microservices applications**, and **any IT component**, including runtimes, databases, middleware and cloud instances.

Custom **optimization goals & constraints** allow you to set the desired **performance, resilience and cost objectives** for each application to be optimized.

AKAMAS

Akamas benefits for Kubernetes apps

Akamas autonomous optimization guarantees the **best levels of performance & resilience** while also ensuring the **best cost efficiency** of your Kubernetes microservices applications, thus avoiding any resource overprovisioning and unnecessary infrastructure/cloud costs.

With Akamas, **developers, performance engineering & SREs** are freed from manual tuning activities and can focus on innovating and operating Kubernetes applications, by automatically receiving actionable insights on potential tradeoffs and recommended configurations.

This results into **higher operational efficiency, business agility** and **competitive advantages**.

-70% decrease in response time with lower fluctuations and peaks

-60% decrease in infrastructure/cloud costs with the same application performance

-80% savings in engineering time spent for manual tuning tasks

AKAMAS

AKAMAS

info@akamas.io