

AKAMAS

SUCCESS STORY

Leverage AI to cut cloud costs and improve your business agility

Akamas AI-powered optimization provides actionable insights on how to safely reduce costs and deliver better services

NAVAN

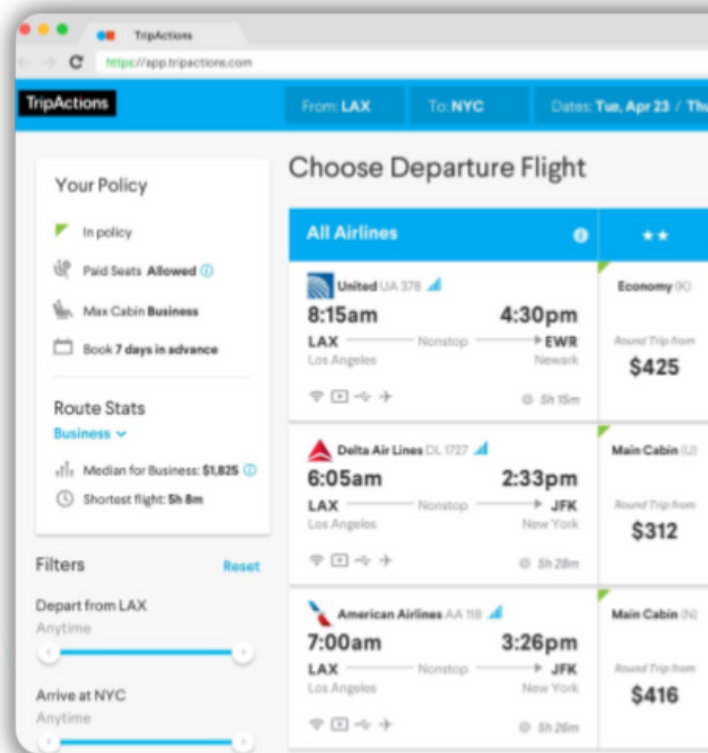


Worldwide leader in travel, corporate card and expense management for businesses

Navan is a leading worldwide travel management, corporate card and expense management company for businesses.

Navan helps today's businesses keep employees informed and safe by taking the right decisions thanks to its modern, user-friendly technology and comprehensive support provided by experienced, knowledgeable travel agents.

Navan enables its users to easily book trips, file receipts, modify reservations, track travel itineraries, chat with travel support agents, receive real-time notifications, and reconcile expenses.



trusted by
+8,800
companies
globally

2,200
employees
worldwide

+60
countries
presence

The key challenge of online businesses

How to balance end-to-end performance, resilience and cost to support business growth and improve service efficiency

These days, any company delivering critical online services is facing a fundamental challenge: how to effectively support business growth while also improving cost efficiency by best balancing end-to-end performance, resilience, and cost.

The complexity of today's applications, whether monolithic or microservices-based applications, on-premise, whether cloud-based or native, makes this balancing a daunting task even for the most skilled performance tuning experts. To add to this, the number of tunable parameters and cloud options keeps growing. For example, AWS lists more than 300 EC2 instance types "for virtually every workload or business need". While this provides cloud teams with plenty of choices for building the infrastructure to support their cloud-based applications, the impact of these choices on the cloud bill can be quite significant and difficult to predict. Finally, the interplay among parameters set at multiple application and infrastructure layers is often hard to anticipate and sometimes even counterintuitive.

After the heavy impact of COVID-19 on the travel business, **Navan** expected a strong business recovery, but their focus to running costs was higher than ever. Despite their great level of expertise and deep knowledge of their own applications, **Navan** performance and SRE teams were challenged by the request to properly tune their applications running on AWS to both reduce the cloud bill and sustain the expected traffic growth.

This is where **Navan** turned to **Akamas AI-powered optimization** with the challenge to tune their most critical service, also referred to as the "TA Server", a Java-based application running on AWS EC2 instances, which was considered the hardest service to tune, as being a sort of black-box.

"Akamas helped us to rapidly mature on the performance tuning front, by allowing us to find an optimal configuration for our application. This resulted in significant cost savings as well as removing barriers to replatforming."

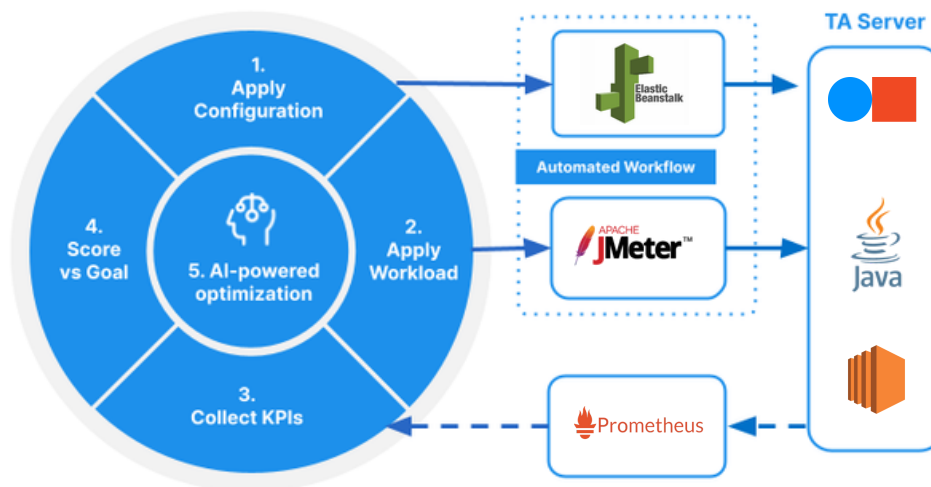
Chris Cholette , VP Productivity and Site Reliability Engineering at Navan

Lower bills with no business impact

Automatically identify the EC2 and JVM configuration that reduces cloud costs and also support higher load

The first **Akamas** optimization campaign at **Navan** was to reduce the cloud bill of their TA Server, while possibly sustaining additional traffic, and also ensuring that both SLOs (e.g. error rate lower than 1%) and internal resource utilization policies (e.g. CPU $\leq 80\%$) were matched.

To analyze this what-if scenario, **Akamas** was deployed into a pre-production, scaled-down environment. All the steps required in each experiment to verify the result of applying each configuration recommended by **Akamas AI** are executed without any human intervention: the TA Server is automatically deployed via Beanstalk, the corresponding JVM parameters and EC2 options are set, and the load mimicking the production load (a mix of API calls for flight, hotel and car-related requests) is applied, also automatically.



High level architecture of the cost-optimization scenario

After fewer than 30 experiments, **Akamas AI** identified an optimal configuration that was capable of providing a significant cost reduction (about 55% less), sustaining almost twice the current load, and also matching both SLOs and resource utilization policies. This result can be achieved by touching several JVM parameters and moving from an r5.4xlarge instance to r5.2xlarge.

Thanks to **Akamas**, **Navan** was able to reduce the cost of its most critical service by choosing the recommended EC2 and JVM configuration.

Take cost reduction to the next level

Adopt a scientific approach to validate serverless compute options before planning migration of your critical services

AWS Fargate is AWS serverless compute offering that allows companies to deploy and manage their applications without having to manage the underlying infrastructure, thus removing any operational overhead associated to EC2 instances.

For **Navan**, moving to Fargate could represent an interesting cost-saving opportunity, unless it would have required a major rework of the application.

Thus, the challenge was how to validate the business case, by verifying the ability to run their critical application on Fargate and avoid any undesired business impact.



Akamas Insights into optimal configurations

Akamas provided the perfect solution to validate the Fargate scenario. This was done by running the TA Server on a EC2 instance type (r5.xlarge) with HW specs similar to the targeted Fargate environment. **Akamas** automatically identified the optimal JVM configuration to run the TA Server on Fargate (and match SLOs).

The positive results of the **Akamas** studies provided the **Navan** team the confidence required to plan the migration of the TA Server to AWS Fargate.

"Tripartions is growing at a massive rate, Akamas helped us create headroom in our most critical application without having to make code changes. In addition, we were able to find a AWS Fargate-friendly configuration, which is a big win for us."

Chris Cholette , VP Productivity and Site Reliability Engineering at Navan

AI-powered Optimization Platform

Akamas is the **AI-powered optimization platform** that enables enterprise and online business to consistently deliver unprecedented level of service performance and resilience at minimum cost. Akamas exploits **patented machine learning** techniques to empower developers, DevOps, performance engineers, SREs and service owners to ensure their services are continuously, automatically and autonomously optimized by identifying the best configuration against any **custom-defined goal and SLOs** for hundreds of parameters at application, middleware, cloud and infrastructure levels.



Full-stack

Any technology, any parameter, any system



AI-driven

Let AI find the optimal configuration



Goal-oriented

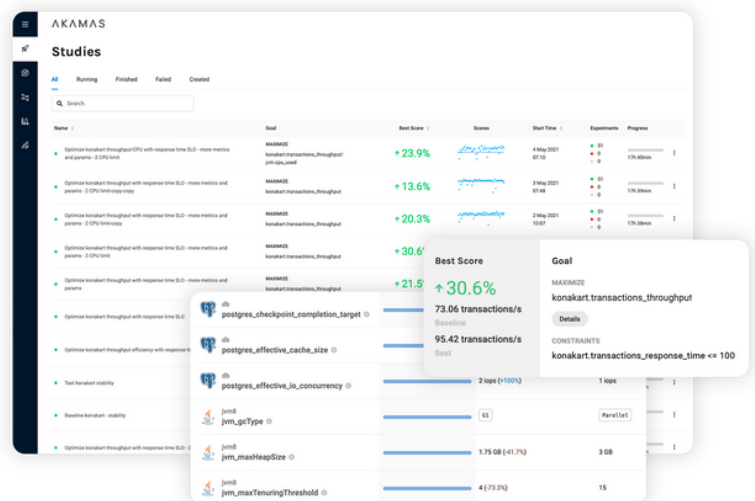
Set your goals & constraints for the right tradeoffs



Closed-loop

Automation that fits your delivery process

Test it out!
Contact us at
info@akamas.io



Milan HQ

Via Schiaffino, 11
20158 MILANO

Boston

211 Congress Street
Boston, MA 02110

Los Angeles

12130 Millennium Drive
Los Angeles, CA 90094

Singapore

5 Temasek Blvd,
Singapore 03898