

AKAMAS

SUCCESS STORY

## **Pioneering innovation in resource optimization with AI: Sisal's journey**

How a global gaming leader leveraged Akamas' AI-powered, full-stack, live optimization platform to safely cut its cloud costs

**Sisal**



## The leading gaming operator since 1946

**Sisal** is a leading gaming operator headquartered in Milan, Italy, directly employing 2,700 people. Since 1946, Sisal has been a leading and responsible operator in the gaming industry, with a diversified portfolio and its omnichannel approach, which allowed Sisal to reach over 35 million customers worldwide.

**Sisal's** strategy rests on three pillars: sustainability, through a constant commitment to the development of the Responsible Gaming program and through the offering of a secure, transparent entertainment model; digital innovation, thanks to its cutting-edge gaming platform oriented towards an omnichannel approach and in-house software and application development expertise to grasp the opportunities offered by the digital transition; and internationalization, with the goal of being awarded contracts for new international concessions on the basis of its solid expertise.

Since 4th August 2022, **Sisal** became part of the Flutter Entertainment group, the world's largest sports betting and gaming operator with a portfolio of globally recognized brands. The group has a diverse portfolio of well-known brands, such as Betfair, Paddy Power, Fanduel, and PokerStars.

With over 1.8 million consumers in the online channel, **Sisal** operates on a fully proprietary technology stack with a strong focus and track record in delivering product innovation and development.

**+75**

Year in business

**47,000**

Points of sales

**35M**

Customers

# The challenge

## Continuously optimize critical Kubernetes microservices for maximum cost savings, with no impact on application performance

**Sisal**, not unlike other large technology-driven operators delivering real-time transactional services, relies on applications and microservices built on Kubernetes, the de-facto operating system of modern cloud-native applications.

Sisal's strategy to run applications on a cloud-native microservices architecture comes with the added complexity of Kubernetes. In particular, Kubernetes resource management poses challenges in efficiently allocating and scaling resources across numerous containers and services.

This approach requires careful planning and monitoring to ensure optimal performance, cost-effectiveness, and resilience. Failing to correctly manage Kubernetes configurations will cause Sisal potential downtime and diminished user experiences. It may force Sisal to systematically over-provision and pay higher cloud costs than it should.

Sisal, operating over 47,000 point of sale (POS) locations, sought to enhance the performance of its in-production microservices-based **Piattaforma Gestione Terminali (PGT)** application. The PGT application is responsible for handling device logins, bet collection, and maintenance tasks for all Sisal POS customers, utilizing connected devices available in each store. PGT is deployed on a Kubernetes cluster and its resource request and response time thresholds are especially stressed at log-in.

The task for Akamas involved implementing **live** (meaning within the production environment) and **full-stack** (simultaneously adjusting multiple technology layers) configuration changes to the application. These adjustments focused on CPU and memory limits within the Kubernetes environments and JVM parameters, such as garbage collector and memory heap size. Throughout the process, it was essential to maintain the key performance indicators (KPIs) identified by Sisal as critical: *throughput and response time*.



# Why Sisal chose Akamas

The IT automation market is flooded with cloud optimization solutions. Some are well-suited to automate traditional performance testing approaches, others focus on individual infrastructure layers such as cloud instances or java configurations.

**Sisal** needed a solutions that could, at the same time:

Support both live optimization of services in production

Support pre-production testing and what-if scenario analysis

Be application-aware to meet high-level competing goals

Actively prevent dangerous configurations through intelligent safety constraints


Traditional performance tuning automation alone isn't sufficient to achieve Sisal's optimization goals. This is where Akamas AI-powered optimization platform comes into play. Akamas' optimization platform offers the distinct advantage of optimizing not only in pre-production environments but also directly within live production settings. This capability allows organizations to fine-tune and improve their applications' performance in real-time without affecting ongoing operations. By utilizing Akamas' technology, businesses can smoothly adapt to shifting demands, optimize resource usage, and uphold exceptional service quality in both pre-production and live production environments.

Akamas' unique attributes - **full-stack, dynamic, context-aware** optimization - draw on live observations, adhere to defined SLOs for secure recommendations, and implement custom safety policies for gradual, incremental changes, ultimately achieving the optimal configuration.

# Achieving a 58% cost reduction in production

Akamas Safety Policy feature allowed risk-free exploration of configuration spaces in search of maximum benefits

The business-critical PGT application manages Sisal devices at betting points and is composed of several microservices, built in Java and deployed on multiple Kubernetes clusters. The microservices are deployed on separate Kubernetes clusters for *Production*, *Business Continuity*, and *Disaster Recovery*.



Total savings result from cumulative pod optimization resizing across farms

Akamas' primary optimization objective involved identifying the optimal application configuration to minimize provisioning costs (right-sizing) while maintaining the customer's service level objectives (SLOs) in terms of **throughput** and **response time**.

Akamas was integrated with existing telemetry and observability tooling built on Dynatrace and started gathering metrics from the system in production.

Upon thoroughly examining the application behavior and taking into consideration the customer's SLO requirements, the Akamas AI-powered optimization engine began proposing potential configurations aimed at achieving the desired cost reduction.

After each configuration was implemented in the production environment, Akamas received feedback and subsequently provided another configuration, ultimately achieving a **reduction in cost of 58%**.

A key feature of Akamas that proved invaluable in this case was its safety policies. Gradual optimization takes into account the maximum recommended change for each parameter (e.g., +/- 5% or customizable) relative to the baseline values.

This approach enables Akamas to progressively explore the configuration space and cautiously approach any potentially risky areas, preventing adverse effects on the system.

*“When contemplating the application of AI-generated configuration changes in a business-critical production environment, we sought a solution that would guarantee our application's stability. Although we had already taken necessary precautions by deploying the app across three different clusters, Akamas' Global Safety feature offered the additional assurance we needed to confidently execute application configuration optimization with peace of mind.”*

**Gabriele Bosisio,**  
Control Room Manager at Sisal

# Multi-year savings through continuous optimization

The Akamas full-stack approach unleashes the power of assistive technology that delivers value immediately and also continues to deliver value over the long term.

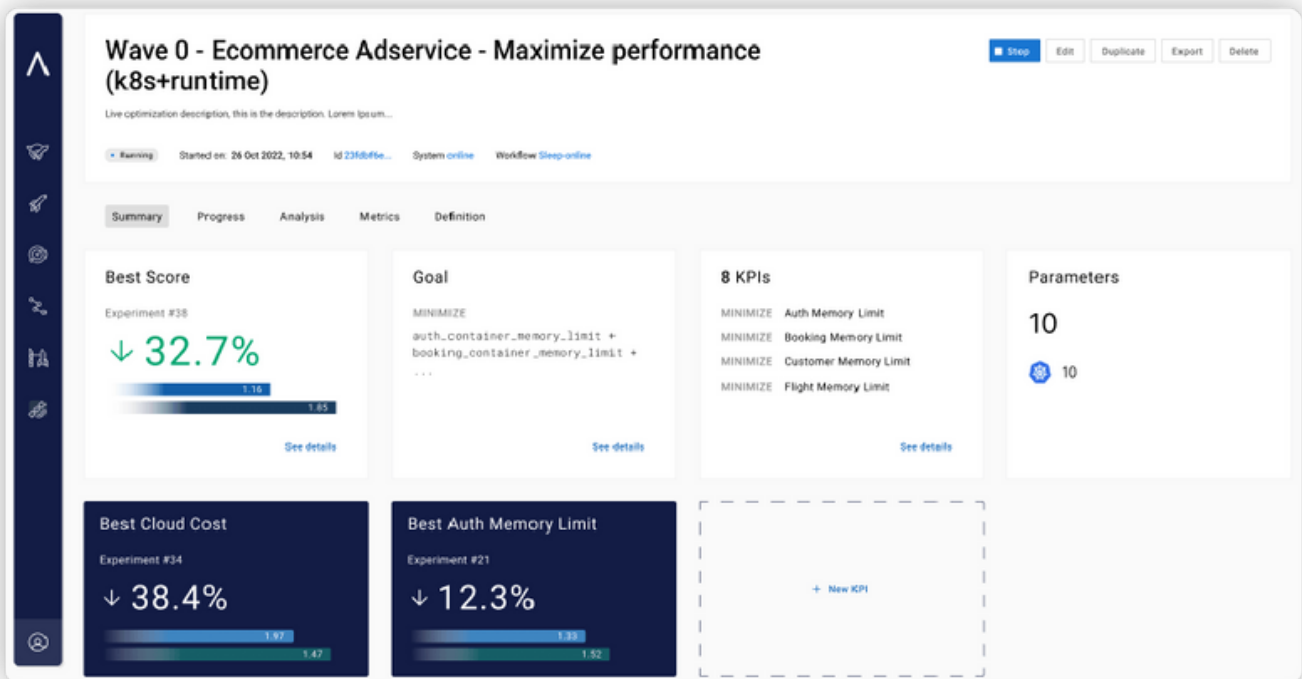
Akamas' full-stack capabilities delivered substantial cost savings through a highly effective optimization process. In less than **two weeks of live optimization**, the Akamas platform showcased the flexibility needed to support simultaneous multi-layer optimizations.

Akamas intervention was not limited to configuring Kubernetes. To successfully achieve the customer's goal of optimizing the Kubernetes layer and reducing the cost function by 58%, the platform simultaneously honed in on JVM memory limits.

The cost reduction was accomplished by concurrently cutting memory by 40% and CPU usage by 68%, addressing the central issue of oversizing associated with JVM memory usage. Akamas' patented AI technology revealed that modifying configurations at the **JVM layer** could also reduce the memory allocated by the Kubernetes pod, leading to remarkable resource optimization and cost savings.

By utilizing live observations and maintaining awareness of the application context, Akamas managed heap size and added a memory safety buffer to mitigate the risk of outages. Subsequently, the platform began optimizing JVM memory, resulting in a 40% reduction in consumption.

Thanks to its full-stack capabilities, Akamas has demonstrated the value of its optimization platform, not just in a one-off, one-layer optimization, but especially as a continuous assistive technology. By achieving significant cost savings in just two weeks, and then continuing to adapt configurations, the platform proves its capacity to support simultaneous multi-layer optimizations and meet evolving customer needs.



Example of Akamas live optimization summary view

*“With over 75 years of experience in the gaming industry, Sisal has consistently been seen as a trailblazer. Microservices applications have become the cornerstone of all modern cloud-native architectures, and it's crucial for us to stay up-to-date with technological advancements. When we discovered the limitless potential of a full-stack, application-aware optimization platform like Akamas, we felt we were right at the forefront of this evolution.”*

**Carlo Monaco,**  
IT Operation Director at Sisal



# Autonomus Performance Optimization

**Akamas** is the AI-powered autonomous optimization solution that enables enterprises and online businesses to deliver unprecedented levels of service performance and resilience at minimum cost. Built by veterans in performance engineering and data science, Akamas exploits advanced machine learning techniques to optimize hundreds of interdependent service configuration parameters while matching both technical and business goals. Akamas customers include leading enterprise organizations in financial and online services.

### Full Stack

any application, any middleware, any database, any cloud, any system

### Smart

patented AI quickly identifying optimal configurations beyond any manual tuning

### Goal-oriented

custom-defined goals translating SLOs and business & technical constraints

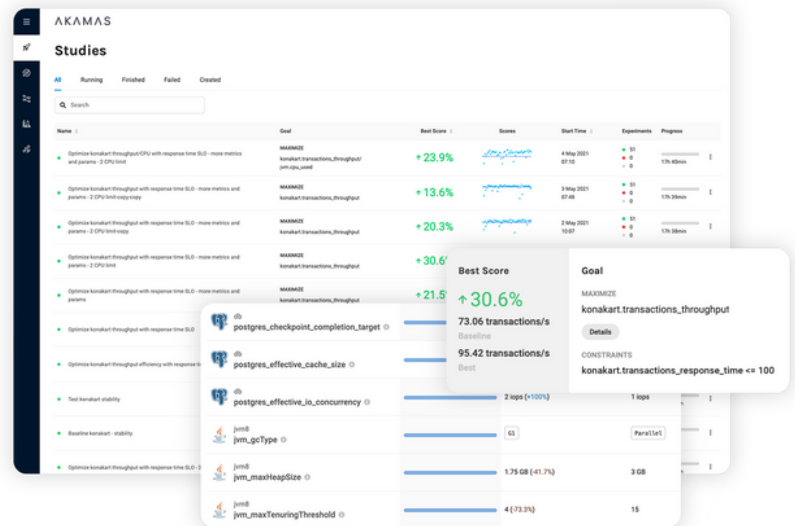
### Automated

custom workflows automating parameter changes, load testing & telemetry collection

### Safe

recommended configurations safely applied to production environments

**Test it out!**  
**Contact us at**  
[info@akamas.io](mailto: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