

Digital Ad Platform for FinServ Saves \$900K+ Annually with Pepperdata on Amazon EKS

A Data-Driven Ad Platform Seeking Cost Mitigation Solutions for Amazon EKS

A large, data-driven, digital advertising platform that partners with banks to operate customer rewards programs was in the process of migrating its significant Spark workloads to Amazon EKS to benefit from the flexibility, control, simplified operational overhead, and reduced costs of a managed Kubernetes service.

The company had been running Apache Spark on Amazon EMR to process a high volume of anonymized and aggregated consumer purchase and other behavioral data, but started migrating these workloads to Amazon EKS to further reduce costs.

As part of its efforts to minimize infrastructure costs, the company adopted [Karpenter](#) to optimize resource utilization by automating node lifecycle provisioning and autoscaling clusters.

However, the company was keen to further increase its margins on Amazon EKS, and sought out additional cost reduction solutions for its newly migrated workloads.

Pepperdata: Autonomous Cost Optimization for Kubernetes

Following migration to Amazon EKS and deployment of Karpenter, the client chose to deploy [Pepperdata Capacity Optimizer](#) to further improve its costs.

Using patented algorithms, Pepperdata Capacity Optimizer autonomously optimized cluster CPU and memory in real time to deliver an average of 30% *additional* cost savings from its Amazon EKS cluster resources.

Capacity Optimizer automatically analyzed resource usage in real time to identify pods where more work can be done. It then communicated this insight to the system scheduler, which added tasks to pods with available resources and spun up new instances only when existing instances were fully utilized.

The result: CPU and memory were autonomously and continuously optimized to reduce cost without the need for application changes or recommendations to be applied manually, safely eliminating the need for ongoing manual tuning.

Pepperdata paid for itself, immediately decreasing instance hours and waste, increasing Spark utilization levels, and freeing developers from manually tuning old apps to focus on building new apps for business growth and innovation.



Realizing Savings Immediately in Both Development and Production Environments

The ad platform company achieved its Proof of Value (POV) with Pepperdata in two phases: first in the company's development cluster to quickly validate the solution robustness and cost saving results, and then in the company's production cluster.

Capacity Optimizer immediately delivered a **22% reduction** of instance hours (vCPU hours) for the company's Spark workloads on their Amazon EKS development cluster. In the second phase of the POV, the cost savings continued in the production environment. By decreasing the number of vCore hours required to run workloads on Amazon EKS, the company realized increased price/performance **equivalent to a monthly cost reduction of approximately \$75K** in their production clusters.

The company was even able to benchmark its savings by temporarily disabling Capacity Optimizer and observing the resulting 22% increase in costs.

Pepperdata Capacity Optimizer Enables Additional Savings No Other Solution Can Provide


While the migration to Amazon EKS brought additional efficiencies, and Karpenter provided cost optimization through rightsizing compute resources, these solutions were not designed to address the problem of waste within applications due to overprovisioning. Spark is well known to be inherently resource-intensive in its design, especially with regard to memory. Even with the most efficient Karpenter-optimized infrastructure in the world, overprovisioned Spark applications will use infrastructure resources inefficiently. **Capacity Optimizer uniquely addresses this inefficiency**, providing the company a complementary benefit to their use of Karpenter.

Pepperdata Cost Optimization at the Application Level

REDUCES INSTANCE HOURS AND COST BY 30% ON AVERAGE	• Autonomous & continuous • No modeling required • Delivers real-time and ongoing savings • Increases utilization
INSTALLS QUICKLY AND DEPLOYS SECURELY	• Installs in <1 hour via Helm chart in Kubernetes • Enterprise-grade secure (SOC 2 Type 2 Compliance) • Delivers visibility into performance and savings in <1 day
EASY TO MAINTAIN → SAVES DEVELOPER TIME	• Requires no application code changes • No recommendations to implement • Eliminates the hassle of manual tuning • Maintains your environment at peak performance


Figure 1: Pepperdata Capacity Optimizer automatically optimizes cluster CPU and memory resources to reduce instance hours and cost, solving the in-application waste problem that no other solution, including Karpenter, can address.

Pods running suboptimally WITHOUT Pepperdata



Without Pepperdata, pods may run at a fraction of their potential utilization. The scheduler is only aware of requested resources and is not aware of what is actually being used, and therefore it cannot pack additional workloads onto pods that appear to be fully utilized.

Pods running at greatest capacity WITH Pepperdata



With Pepperdata, the blinders from the scheduler are removed, enabling it to pack pods according to their actual utilization. The scheduler can make full use of existing pods, increasing capacity and throughput and reducing cost without having to spin up new instances.

Figure 2: Pepperdata enables pods to run at the greatest capacity and highest efficiency on Kubernetes clusters.

Reinvesting Savings for Continued Growth

Pepperdata Capacity Optimizer helped this digital ad platform company achieve significant efficiencies and cost savings through powerful, real-time, and continuous optimization. Working safely and autonomously in the background, Capacity Optimizer minimized in-application waste to increase utilization levels while also eliminating the hassle and uncertainty of manual application tuning, delivering this customer significant and measurable ROI.

The company can now reinvest their ongoing savings from Capacity Optimizer in continued migration of workloads to EKS, as well as in further business innovation and expansion.

About Pepperdata

Pepperdata cost optimization for data-intensive workloads such as Apache Spark is the only solution that autonomously reduces overprovisioning waste to deliver an average 30% cost savings continuously and in real time with no application changes or manual tuning. Deployed on over 20,000+ clusters, Pepperdata Capacity Optimizer optimizes Kubernetes and YARN resources in real time in some of the largest and most complex environments in the world. Founded in 2012, Pepperdata has helped top enterprises such as Citibank, Autodesk, Magnite, Royal Bank of Canada, members of the Fortune 10, and mid-sized companies save over \$250 million. For more information, visit [pepperdata.com](https://www.pepperdata.com).



Pepperdata, Inc.
530 Lakeside Drive
Suite 170
Sunnyvale, CA 94085

 **Start a Free PoV**
www.pepperdata.com

 **Send an Email**
eval@pepperdata.com