



Pepperdata for Amazon EKS

Reduce your cost of running Spark, microservices, and Amazon EMR on Amazon EKS—autonomously, continuously, and in real time—with zero code changes and no manual intervention.

Benchmarking on Amazon EKS

Pepperdata selected [TPC-DS](#), the decision support benchmark from the Transaction Processing Performance Council, as the standardized workload model for its benchmarking using a GitHub repository that closely adheres to the TPC-DS model. Pepperdata ran this benchmark “out of the box” to create an unofficial audited benchmark as defined by TPC.

Using a 1 TB dataset on 500 nodes running 10 parallel applications (with 275 executors per application) to simulate a multi-tenant environment, Pepperdata found that Capacity Optimizer:

- **Reduced cost:** Reduced instance hours consumed by 41.8%
- **Increased throughput:** Reduced total workload run time by 45.5%

This benchmark showed the value of Pepperdata for batch workloads running on Amazon EKS at scale. [Read more here.](#)

Special, Limited Time Offer

Get started today with a free, 2-day waste assessment of your Amazon EKS environment. Visit pepperdata.com or contact us at sales@pepperdata.com.

Driving Savings for Global Enterprises



Pepperdata for Amazon EMR on Amazon EKS

Kubernetes has become the de facto standard for managing applications and services in the cloud. Kubernetes automates and simplifies your container workflows while providing the scalability to adjust seamlessly to ongoing changes in traffic. Many organizations choose to run their Kubernetes workloads on Amazon Elastic Kubernetes Service (Amazon EKS) to gain the additional deployment and scaling benefits offered by a fully managed service. Amazon EKS enables a greater degree of container or node management by automating load distribution and parallel processing and making it easy to run tooling and plug-ins from the Kubernetes open-source community.

Savings at Scale

According to a [recent survey of enterprise users](#), Apache Spark is one of the primary workloads being deployed on Kubernetes. This same survey identified “significant or unexpected spend” as the top challenge to Kubernetes adoption. Pepperdata Capacity Optimizer Next Gen, a cost-optimization software solution battle-tested in some of the world’s largest and most complex computing environments, delivers automatic cost control for Kubernetes—including Spark on Kubernetes—without the need for manual application tuning. Capacity Optimizer Next Gen empowers you to accelerate your savings on Amazon EKS and increase the overall value and effectiveness of your cloud investment.

Real-Time Capacity Optimization

Developers typically configure their pods to request more resources than needed with the goal of ensuring that their applications run to completion. While this practice may be manageable in development environments, it can create expensive waste when jobs are moved into production. **Pepperdata Capacity Optimizer Next Gen** offers a patented **Continuous Intelligent Tuning that autonomously reduces costs and optimizes resource utilization in real time** for Spark and microservices on Amazon EKS and also for Amazon EMR on EKS.

For batch workloads such as Spark, Capacity Optimizer Next Gen is able to improve node efficiency by enabling the YARN scheduler to launch tasks based on hardware utilization rather than relying on allocations, which by design contain waste in the form of overhead.

Capacity Optimizer Next Gen further enables the autoscaler to scale only based on actual usage instead of allocations, which ensures that existing nodes in the cluster are fully utilized before the autoscaler adds additional nodes. And it does all of this in real time in response to dynamically changing workloads and datasets. As a result, Capacity Optimizer Next Gen always maintains nodes in a sweet spot of optimal utilization, which cannot be done manually. In this way, your applications consume only the resources they actually use, which translates into significant cost savings. Capacity Optimizer Next Gen ensures that you pay only for the resources you actually use, rather than paying for higher levels of allocated resources that go unused and wasted.

Significant Cost Reductions and Performance Enhancements at Scale


Pepperdata Capacity Optimizer helps improve the performance and cost of Spark workloads running on Amazon EKS by automatically packing additional pending pods onto underutilized nodes, increasing node utilization and reducing the need for additional nodes. According to a recent benchmarking study, Capacity Optimizer enabled a standard Spark workload running at scale on 500 Amazon EKS nodes to run at 41.8 percent reduced cost, as measured by reduced instance hour consumption. In addition to delivering significant savings on Amazon EKS even at scale, in this same study, Capacity Optimizer also enabled the Spark workload to run 45.5 percent faster, as measured by decreased run time.

Advanced Visibility and Optimization Flexibility

In addition to providing immediate autonomous improvements in cost and performance, Pepperdata dashboards provide advanced visibility into Spark applications and the clusters in which they are running, including Spark metrics, containers, clusters, pods, nodes, and workflows, as well as showing the total realized savings and potential savings available at the cluster level. This advanced visibility translates into greater understanding of your workloads and the ability to adjust optimization levels to your desired price/performance SLAs.

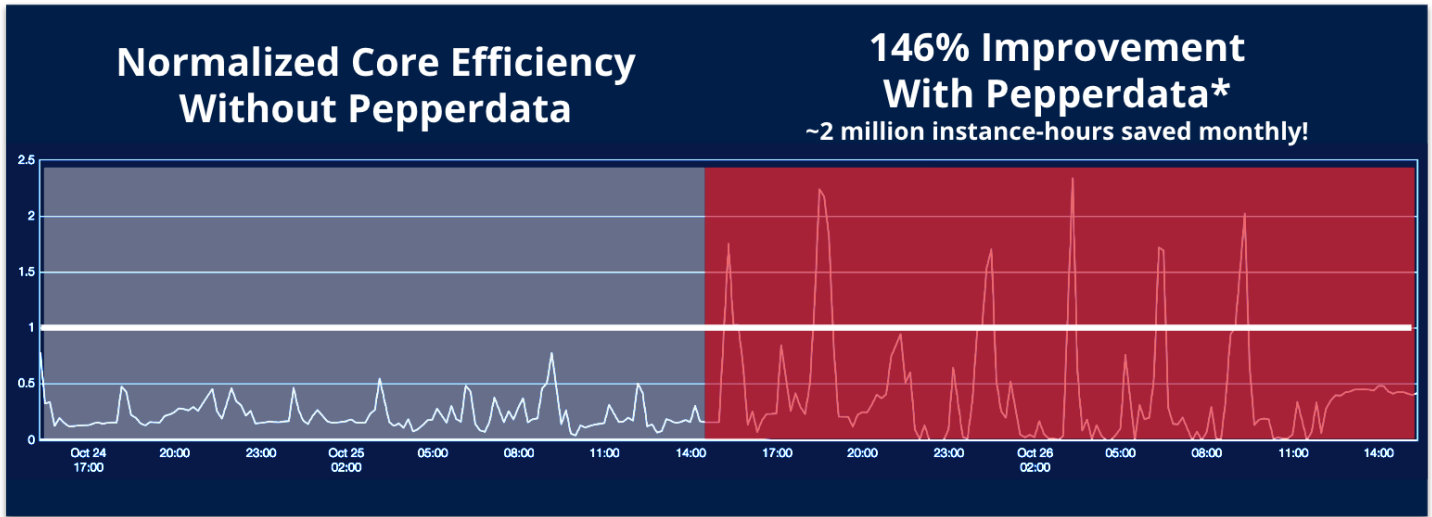
For microservices workloads, Capacity Optimizer Next Gen acts like a turbocharged Vertical Pod Autoscaler, automatically adjusting the CPU and memory requests of your pods to the actual usage of your services. Capacity Optimizer Next Gen aligns the pod resource requests with actual usage, reducing resource requests to better match the specific usage levels of your applications and thus minimizing wasted resources and costs.

For Amazon EMR workloads running on Amazon EKS, just as with general Spark workloads, Capacity Optimizer Next Gen delivers additional performance improvements and significant cost reductions. Initial informal data from a current Pepperdata customer shows a 42.5 percent reduction in instance hours when Pepperdata Capacity Optimizer with Autoscaling Optimization is enabled. Autoscaling Optimization intelligently augments the native autoscaler to ensure all pods are fully utilized before additional pods are launched, eliminating waste and reducing costs.



The same customer running Amazon EMR on Amazon EKS experienced a 146 percent Improvement in efficiency once it began running Capacity Optimizer with Autoscaling Optimization, as shown in the figure below.

Driving Performance Uplifts for Amazon EMR on Amazon EKS



* Actual customer data from October 2023

Capacity Optimizer's ability to autonomously optimize high-performance workloads such as Amazon EMR on Amazon EKS helps ensure that cloud resources are used efficiently and cost effectively and thus makes the cloud an even more attractive and viable option for such workloads.

In short, Pepperdata Capacity Optimizer has been shown to improve the performance and reduce the cost of running both Apache Spark and microservices workloads as well as Amazon EMR on Amazon EKS. Pepperdata Capacity Optimizer delivers these savings at scale through its Continuous Intelligent Tuning that eliminates wasted resources from every app across your entire cluster in real time. By reducing instance hours continuously and autonomously throughout your ever-changing Amazon EKS environment, Pepperdata Capacity Optimizer reduces costs and keeps your clusters running at the sweet spot of peak utilization..

Pepperdata: Your Trusted Partner

Pepperdata is the only cost optimization solution that delivers significant cost savings continuously and in real time on Amazon EMR and Amazon EKS with no application changes or manual tuning. Our customers include the largest and most complex clusters in the world, at top enterprises such as Citibank, T-Mobile, Autodesk, Securonix, Royal Bank of Canada, and those in the Fortune 5.

Supported Technologies

- Amazon EKS
- Amazon EMR for EKS
- Apache Spark

About Pepperdata

Pepperdata is the only cost optimization solution that delivers up to 47% greater cost savings—continuously and in real-time—on Amazon EMR and EKS with no application changes or manual tuning. Our customers include the largest, most complex, and highly-scaled clusters in the world, at top enterprises such as Citibank, T-Mobile, Autodesk, Securonix, Royal Bank of Canada, and those in the Fortune 5. 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