

Pepperdata for Amazon EKS

Slash the cost of running Spark, microservices, or Amazon EMR on Amazon EKS by up to 47%—autonomously, continuously and in real time—with zero code changes and no manual intervention.

Driving Growth for Global Enterprises



TPC-DS Benchmarking

TPC-DS is the Decision Support framework from the Transaction Processing Performance Council. TPC-DS is an industry-standard big data analytics benchmark. Pepperdata's work is an unofficial benchmark as defined by TPC.

Using a 1 TB dataset on 30 nodes with 225 executors and 103 TPC-DS jobs, Pepperdata found that Capacity Optimizer Next Gen:

- **Improved performance**
Decreased query duration by 30%
- **Increased throughput**
Increased workload capacity by 42%

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 rackspace@pepperdata.com.

Supported Technologies

- Amazon EKS
- Amazon EMR for EKS
- Apache Spark

Kubernetes has become the standard for managing modern-day applications and services in the cloud. But despite its power and flexibility, running Kubernetes at scale can pose some challenges—most notably, controlling costs without constant manual application tweaking.

Real-Time Capacity Optimization

Developers typically configure their pods to request more resources than needed. While this practice may be manageable in development environments, it can create needless waste when jobs are moved into production. Pepperdata Capacity Optimizer Next Gen slashes cost and optimizes resource utilization for Spark and microservices on Amazon EKS and on Amazon EMR on EKS.

- **For Spark applications**, Capacity Optimizer Next Gen packs additional pending pods onto underutilized nodes, increasing node utilization and reducing the need for additional nodes, which translates directly to reduced costs.
- **For microservices**, Capacity Optimizer Next Gen acts like a vertical pod autoscaler and works with the horizontal pod autoscaler to align pod resource requests with actual usage so that pods, workloads, and nodes can be scaled efficiently.

Pepperdata Capacity Optimizer Next Gen works autonomously in the background in real time, maintaining your clusters continuously at the sweet spot of optimal utilization. This eliminates the hassle of manual tuning and frees your developers to focus on production and innovation.

Optimized Autoscaling

While Amazon EKS autoscaling provides the elasticity required for Spark applications, an improperly configured autoscaler can lead to increased costs. Pepperdata further optimizes your Amazon EKS environment by dynamically tuning your autoscaler based on *actual* and rather than *requested* resources. By packing more pods onto existing nodes, Pepperdata automatically reduces the number of nodes needed, allowing you to achieve up to 30% additional cost savings with no code changes.

Pepperdata: Your Trusted Partner

Pepperdata is the only cost optimization solution that delivers up to 47% 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.