

Capacity Optimizer

Real-Time Cost Optimization Without Recommendations

Reduce workload costs by up to 47% without application changes

Savings at Scale Through Continuous Intelligent Tuning

Pepperdata Capacity Optimizer continuously and autonomously reduces the waste and cost in your clusters by increasing node level utilization in real time without the need for application changes. We call this **Continuous Intelligent Tuning**.

Node Level Optimization

Pepperdata's Continuous Intelligent Tuning enables the scheduler to make use of the node resources that are allocated but not used. Capacity Optimizer makes the allocated but unused resources visible to the scheduler so that it can launch more workloads on existing nodes instead of adding new nodes.

Autoscaling Optimization

Pepperdata also optimizes autoscaling by making sure that new instances are launched only when the existing instances are fully utilized.

The result: CPU and memory are autonomously optimized to run more workloads to increase your savings.

Reduce Workload Costs with Autonomous, Real-Time Optimization

Controlling costs is incredibly difficult. Typically, teams will optimize cost at the platform or infrastructure level with instance right sizing, along with the implementation of Reserved Instances and Savings Plans. While these modifications save money, there is still waste in the application itself.

Pepperdata works autonomously in real time across the entire cluster to increase virtual capacity of every node, thus enabling the scheduler to reclaim what would be wasted application resources, or allocations that go unused.

"No one outside of the operations or finance teams [now] has to think about cost management for reporting, and that... is immensely beneficial to the business."

—Ben Smith, VP of Technical Operations, Extole

Nodes running suboptimally without Pepperdata

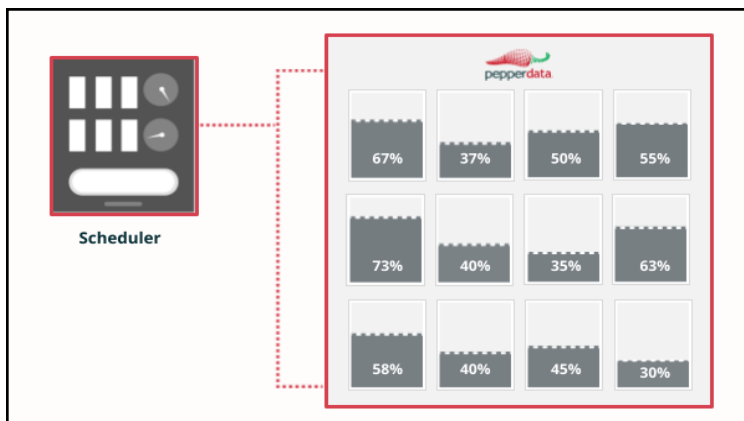


Figure 1: Without Pepperdata, the scheduler does not see that allocated node resources are not fully utilized, leading to wasted capacity and cost.

Nodes running at greatest capacity with Pepperdata




Figure 2: Pepperdata increases capacity, utilization, and savings by communicating to the scheduler that nodes still have allocated but unused resources, allowing the scheduler to add more jobs to existing nodes.

Benefits of Capacity Optimizer

Pepperdata Capacity Optimizer eliminates the need to change applications based on recommendations by working immediately in real time to keep your instances at their optimal utilization with no code changes. The direct benefits to you include:

1. Decreased Instance Hour Consumption/ Reduced Costs

By maintaining your clusters in their sweet spot of utilization, Capacity Optimizer reduces hardware usage by up to 47 percent. The decreased instance hours translate directly to reduced cost and a lower monthly bill.

2. No Manual Tuning, No Recommendations, No Application Changes

Because Capacity Optimizer works continuously in real time, no application changes or manual interventions are ever necessary. You'll never have to implement a recommendation; Pepperdata does all the work autonomously. This means greater savings and more engineering hours freed up for innovation and higher value activities.

"On average, we were saving 50% on our costs because of Pepperdata."

—Mark Kidwell, Chief Data Architect of Data Platforms and Services, Autodesk

Seamless Installation

Capacity Optimizer installs in under an hour on the cluster(s) you choose through either a bootstrap script for Amazon EMR clusters or with a Helm chart for Amazon EKS environments.

Supported Technologies

- Amazon EMR on EC2
- Amazon EMR on EKS
- Apache Spark on Amazon EKS



Pepperdata and the Pepperdata logo are trademarks or registered trademarks of Pepperdata Inc. All other trademarks are the property of their respective owners. Pepperdata reserves the right to change this document without notice. To ensure you have the latest version of this document, visit www.pepperdata.com.

Cost Savings and Efficiency Improvement for Amazon EMR on Amazon EKS




Figure 3: 42.5 percent monthly instance hours saved

For Amazon EMR on EKS customers, Capacity Optimizer intelligently augments the native autoscaler to ensure all pods are fully utilized before additional pods are launched.

A Pepperdata customer recently achieved a 42.5 percent savings of instance hours for a monthly estimated savings of over \$62K while implementing Pepperdata Capacity Optimizer on their Amazon EMR on EKS cluster.

Actual Customer Estimated Monthly Savings

	Normalized Instance Hours Saved	r5.8xlarge Instance Hours Saved	Cost Savings*
Customer's Staging Environment	101,221	1,581	\$3,187.30
Customer's Production Environment	~2 million	31,173	\$62,844.77

*Assuming an on-demand price of \$2.016/hour for an r5.8xlarge instance

Figure 4: Production Environment \$62K monthly savings

Free Savings Assessment

Wondering how many wasted Instance Hours you might be paying for? [Sign up for a free, two-day Pepperdata Savings Assessment](#) to see how much you can save.

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 Amazon 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.