pepperdata

Capacity Optimizer Classic

Do More With Less

Extract the most from your cloud or on-premises investment and reduce big data costs by up to 47%

Maximize Resource Utilization and Improve Cluster Efficiency

You've made a significant investment in your on-prem or cloud infrastructure, but how do you know you're extracting the maximum value out of it? Pepperdata Capacity Optimizer Classic provides you that peace of mind.

Capacity Optimizer Classic autonomously optimizes your clusters continuously and in real time, reduces costs by eliminating waste and maximizing resource utilization, and eliminates the need for tweaking and tuning your clusters and applications. It does this by analyzing the resource usage of each of your cluster's nodes in real time and uses machine learning to make thousands of resource allocation decisions per second. It identifies where more work can be done and adds tasks to nodes with available resources. The result: CPU, memory, and I/O resources are autonomously optimized to increase utilization, and waste is eliminated in both Kubernetes and traditional (Hadoop/YARN) big data environments. Even the most experienced operator dedicated to resource management can't make manual configuration changes with the real-time immediacy of Capacity Optimizer Classic.

Gain Deep Insight Under a Single Pane of Glass

Understanding how your applications are performing across your clusters is essential for helping you meet price/performance SLAs. Capacity Optimizer Classic cuts troubleshooting time and helps you maintain those SLAs through three specific features:

Application-Aware Observability

- Visualize thousands of application-level metrics in one aggregated view. View a summary across all your applications or drill down one by one.

Side-by-Side Application Comparison

- Choose any two applications (or different runs of the same application) to compare their performance characteristics for expedited troubleshooting.

Optimization Recommendations

- Minimize inefficiencies and bottlenecks caused by factors such as insufficient parallelism, memory waste, OOM errors, Spark failures, unused GPUs, and read/write skew.

Dynamic Autoscaling

Capacity Optimizer complements traditional EMR autoscaling by reducing resource waste on your cluster before EMR autoscaling is enabled. On top of Amazon EMR, Capacity Optimizer can reduce the number of cores by up to 63%, active nodes by up to 67%, and CPU idle time by up to 30%.

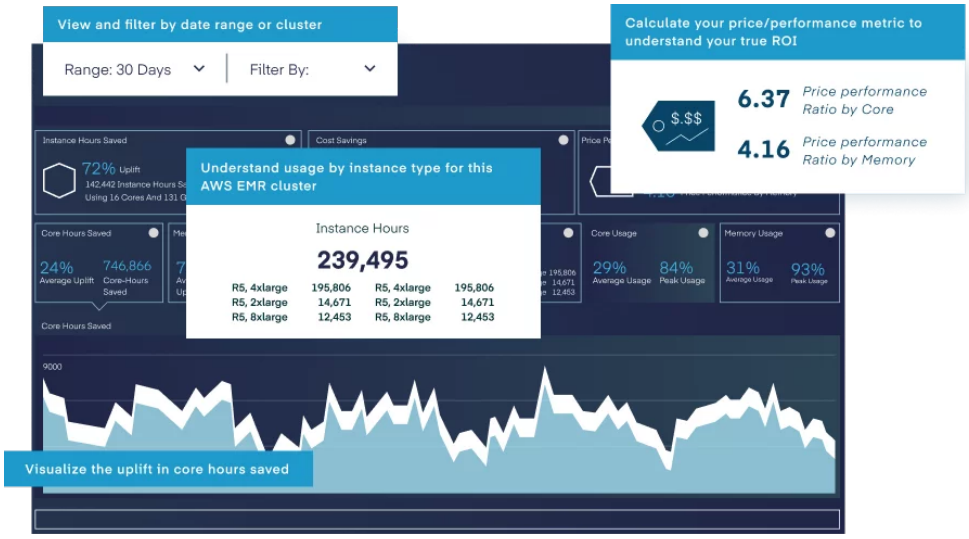

You're In Control

- Install Capacity Optimizer on the cluster(s) you choose and experience immediate resource and cost reductions.
- Relax with the peace of mind that your clusters are operating at the highest levels of resource and cost efficiency.
- Automatically extend the life of your on-prem hardware investments by running up to 47 percent more workloads on your existing system.
- Free your developers from the drudgery of manual tweaking and empower them to focus on innovation and growth.

Supported Technologies

- Amazon EKS
- Amazon EMR for EKS
- Apache Spark
- Google Cloud Platform (GCP)
- Microsoft Azure

Reclaim Reserved but Unused Resources

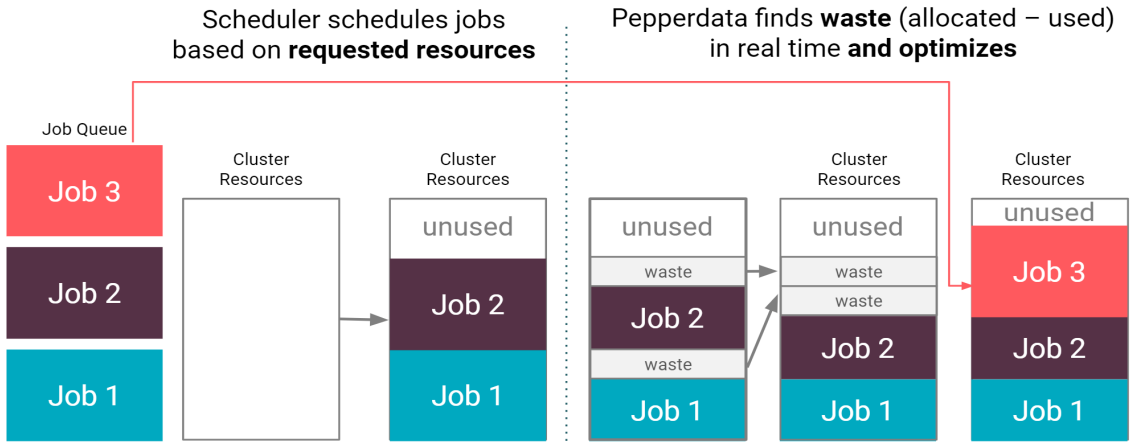

"Capacity Optimizer automatically tunes cluster capacity for workloads in real time, avoiding lots of manual tuning and helping to save money."

—Chief Data Architect, Fortune 1000 Company

Ready to get started with a free Pepperdata Waste Assessment?

Reach out to us at rackspace@pepperdata.com or visit www.pepperdata.com/free-waste-assessment to see how much application waste is in your environment.

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.