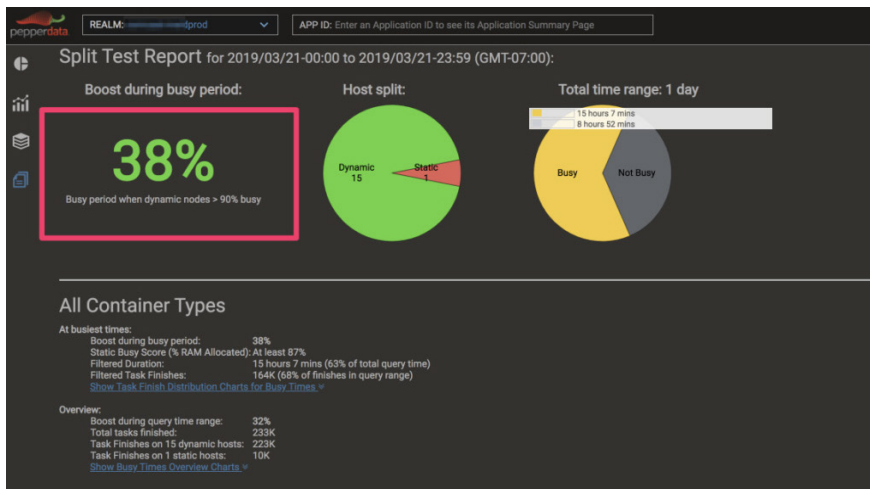


# Capacity Optimizer

Automatically optimize your big data environments at scale.



Capacity Optimizer Split Test

## Get Up to 3X Price-Performance Improvement with Autonomous Optimization

Pepperdata® Capacity Optimizer automatically optimizes your cluster resources, recapturing wasted capacity so you can run more applications and get the most out of your infrastructure investment.

Capacity Optimizer enables you to:

- Get up to 3X price-performance improvement on top of AWS autoscaling.
- Recapture wasted capacity, run more applications, and get the best ROI from your infrastructure investment.
- Optimize each host's ability to run an optimal number of containers and to run the same number of workloads on fewer instances.

## Immediately Improve Big Data Cluster Throughput

On a typical cluster, Capacity Optimizer uses machine learning (ML) to make thousands of decisions per second, analyzing the resource usage of each host in real time. The result: CPU, memory, and I/O resources are automatically optimized to increase utilization, and waste is eliminated in both Kubernetes and traditional (YARN) big data environments.

Capacity Optimizer identifies where more work can be done and adds tasks to nodes with available resources. Even the most experienced operator dedicated to resource management can't make manual configuration changes with that precision and speed.

## Boost Autoscaling Performance and Reduce Your Cloud Costs

In cloud environments, autoscaling provides the elasticity that you need for your big data workloads, but it often leads to uncontrolled costs. Cloud providers provision infrastructure based on the peak needs of workloads. This guarantees that maximums are met but can create a lot of provisioning waste—the very waste that Capacity Optimizer identifies and returns to you in the form of optimized, available resources to run more jobs.

Whatever your cloud platform, Capacity Optimizer uses autonomous optimization to intelligently augment autoscaling and ensure that all hosts are fully utilized before additional hosts are created. The net effect is that horizontal scaling is optimized and waste is eliminated.

**“We achieved a 30% compute boost & continuous optimization. With Pepperdata, we immediately had a major win. They saved us millions on unnecessary hardware expenditures.”**

Technology Director, Leading Global Travel Company

### Pepperdata Experience

- We continuously monitor over 275 customer production clusters.
- These clusters have over 50,000 hosts spanning all big data distributions and hardware configurations.
- We monitor more than one billion jobs and collect over 900 trillion data points annually.

### Pepperdata Expertise

- We have extensive experience with Fortune 1000 customers.
- Our vast real-time proprietary data informs best practices & subsequent platform tuning.
- Our extensive experience with Fortune 1000 clusters provides insight to help you select the best hardware and software stack for your applications.
- Our experts help you squeeze the most performance out of your existing capacity & accurately forecast capacity needs.

### Pepperdata Support

- Amazon Web Services
- Azure
- Google Cloud Platform
- Qubole
- Spark
- HDFS
- YARN
- MapReduce
- Kafka
- Kubernetes
- Tez
- Impala
- Hive
- HBASE

## About Pepperdata

Pepperdata provides analytics performance management that guarantees SLAs and reliability for your infrastructure and apps.

©2021 Pepperdata Inc. All rights reserved. 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](http://www.pepperdata.com).



3945 Freedom Circle, Suite 920  
Santa Clara, CA 95054

**Start a Free Trial**

[www.pepperdata.com/trial](http://www.pepperdata.com/trial)

**Send an Email**

[eval@pepperdata.com](mailto:eval@pepperdata.com)