



# Top Online Bank Realizes 30-75% Cost Reductions on Amazon EMR, Invests Savings in Further Expansion

## About The Client:

The client is a leading online US-based bank with a significant Amazon EMR footprint encompassing dozens of servers processing 150 billion cells of data.

## Challenge:

The online bank had to achieve a corporate initiative to optimize and reduce its rising Amazon EMR cost while maintaining defined SLAs for its data-intensive workloads.

## Solution:

The company used Pepperdata Capacity Optimizer for real-time, automated resource optimization to improve the utilization of CPU and memory of its data workloads for reduced costs.

## Results:

Once installed on seven of the bank's Amazon EMR clusters, **Pepperdata Capacity Optimizer immediately demonstrated resource utilization improvements for cost savings of 30-75 percent per cluster**—enabling the bank to achieve its corporate cost cutting mandate.

The bank is **now realizing an average cost savings of \$140,000 per month** across its Capacity Optimizer-enabled clusters.



## A Leading Online Bank with a Cost-Cutting Mandate

For nearly two decades, a large online bank has been providing its millions of customers across the US with a broad range of financial services.

Underpinning these efforts are the bank's advanced credit decisioning and machine-learning models that are based on more than 150 billion cells of data.

With this significant data footprint, the bank is a major Amazon EMR customer and was running a mix of Apache Spark and Apache Hive workloads—25 percent Hive and approximately 75 percent Spark.

After consistently seeing its Amazon EMR spend increase month to month, the bank instituted a corporate initiative to increase its resource utilization and reduce the hosting cost of its data workloads—all while maintaining performance SLAs.

As part of this initiative, the bank planned to migrate all of its Hive workloads to Spark with the subsequent goal of eventually moving the Spark applications to Amazon EKS.

## Seeking Solutions to Reduce Resource Spending on Amazon EMR

Given the bank's corporate savings mandate, its executives were receptive to a variety of cost optimization options, especially those that could be implemented quickly and without requiring development teams to build custom optimization software.

The bank was introduced to Pepperdata by its AWS Senior Account Manager, who had previously observed Pepperdata's immediate capabilities of increasing workload utilization of CPU and memory to deliver an average cost savings of 30 percent. As part of this evaluation, the bank's large Amazon EMR environment was identified as an area of potential savings.

Transacting with Pepperdata through the AWS Marketplace for a free proof of value (POV), the bank installed Pepperdata on twenty-two of its development and production Amazon EMR on EC2 clusters. Pepperdata Capacity Optimizer was then enabled on seven of the clusters for resource optimization. Pepperdata observability features were deployed on the remaining fifteen to provide both cluster- and application-level visibility.

## Pepperdata Capacity Optimizer Delivers Transformative Results

Once installed on the bank's Amazon EMR clusters, Pepperdata Capacity Optimizer went to work without delay. Instance hours were reduced immediately, resulting in a demonstrated savings of 30-75 percent per cluster. The graph underneath shows the effect of enabling Capacity Optimizer on the number of instance hours required to process the bank's workloads.

Pepperdata Capacity Optimizer provided the system scheduler with real-time visibility into actual resource usage of the bank's workload executors, helping the scheduler identify existing capacity so it could add more tasks to executors that were underutilized.

By increasing the resource utilization of running executors, the cluster autoscaler then only spins up new instances when existing ones are fully utilized.

In turn, this real-time, automated resource optimization immediately reduced the number of instance hours needed to run the workloads and reclaimed memory hours without any need for manual tuning, applying recommendations, or changing application code. The engineering teams could now pivot their focus towards revenue-generating projects versus manually tuning these workloads.

**The average cost savings of \$140,000 per month across all Capacity Optimizer-enabled clusters directly helped the bank achieve its corporate savings mandate.** These savings results were realized immediately and automatically, without any code changes to the company's Spark or Hive applications, and without redirecting any engineering resources away from existing activities.

After enabling Pepperdata Capacity Optimizer, a senior manager on the bank's FinOps team observed that the company-wide Amazon EMR spend that had previously been on a steady increase started to decrease for the first time ever in the company's history.




Figure 1: Capacity Optimizer immediately reduced the number of instance hours required to run the bank's workloads.

# Pepperdata Capacity Optimizer Observability Complements Resource Optimization with a Comprehensive Dashboard

On the remaining clusters, Pepperdata's observability dashboards and data have helped the bank's operations team identify and resolve issues faster, better understand the details of their infrastructure, and measure performance against SLAs.

In fact, Pepperdata's observability dashboard even helped the bank identify a critical bug in related software that they are resolving to further improve performance.

Following the successful deployment of Pepperdata Capacity Optimizer, one of the senior implementation leads commented,

***"Pepperdata promised no code changes, and we have not made a single code change. We haven't needed to work with a single developer yet to deploy Pepperdata."***

—Senior Implementation Lead,  
Top Online Bank



## Savings Fuel Future Expansion Plans Across Amazon EMR and Amazon EKS

Based on the successes to date, the bank expects to achieve an automatic 30 percent cost savings on Amazon EMR moving forward.

With Pepperdata Capacity Optimizer's real-time, automated resource optimization keeping its data workloads running at optimal price/performance, the bank plans to further expand its data footprint on Amazon EMR while migrating its Spark workloads to Amazon EKS for improved performance.

The bank's operation team also plans to use Pepperdata's observability dashboards to help inform this migration and measure the efficiency of their new Amazon EKS environment post migration.

### About Pepperdata

Pepperdata resource optimization for data workloads on Kubernetes and Amazon EMR automatically increases utilization levels by up to 80% to reduce overprovisioning waste, delivering an average 30% cost savings—automatically, continuously, and in real time—with no application code changes, recommendations, or manual tuning.

Deployed on over 30,000+ clusters, Pepperdata Capacity Optimizer optimizes Kubernetes resources in some of the largest and most complex environments in the world. Since 2012, Pepperdata has helped companies ranging from startups and mid-sized ISVs to top enterprises such as Citibank, Autodesk, Magnite, Royal Bank of Canada, and members of the Fortune 10 save over \$250 million. For more information, visit [www.pepperdata.com](http://www.pepperdata.com).

Pepperdata, Inc.  
530 Lakeside Drive  
Suite 170  
Sunnyvale, CA 94085



**Start a Free PoV**  
[www.pepperdata.com](http://www.pepperdata.com)



**Send an Email**  
[eval@pepperdata.com](mailto:eval@pepperdata.com)

