

# Financial Services Giant Saves \$20M with Pepperdata Real-Time, Automated Resource Optimization

## About The Client:

A leading Fortune 100 multinational investment bank and financial services corporation headquartered in the United States.

## Challenge:

The company's on-prem/private cloud Hadoop cluster, which provided a data analytics platform to support its global business operations, was experiencing exponential growth. With ingested data doubling year over year, the company needed to add costly compute resources at a rate that quickly became unsustainable.

The associated maintenance, support, climate control, and real estate expenses required for these compute resources only added to the overall cost. The company needed to gain control of runaway IT resource spend while maintaining optimal workload performance levels and meeting business-critical SLAs.

## Solution:

Pepperdata Capacity Optimizer's real-time, automated resource optimization improved the utilization of CPU and memory for the customer's on-prem data workloads to recapture wasted resources and reduce infrastructure cost. Capacity Optimizer also provided full-stack observability into the performance of the client's data analytics platform to maintain business-critical SLAs.

## Results:

**With Pepperdata real-time, automated resource optimization, the financial services giant achieved a 30 percent reduction in annual server costs—seeing approximately \$20 million in savings over 7+ years** without ongoing manual tuning or application code changes.



## FinServ Provider Facing Uncontrollable Big Data Costs

A leading Fortune 100 multinational investment bank and financial services corporation headquartered in the U.S. needed a real-time, automated solution to address the problem of runaway data center growth and constantly increasing spend for its compute infrastructure. The company provides banking services and products for companies, governments, and institutions in over a dozen countries and is a financial linchpin of the global economy.

The company deployed a Cloudera platform running Hadoop and Apache Spark with a global footprint of approximately 17+ petabytes of data. The platform used the Hadoop cluster as a data lake for storage. The company then managed, processed, and analyzed this data in the cluster using artificial intelligence (AI) and machine learning (ML).

This valuable data was the foundation for many of the company's decision-making and business development processes, and was used to generate equity and fixed-income research as well as economic, geographic, and product-specific analyses.

The company initially engaged Pepperdata in 2017 when the I&O team was considering resource optimization options to improve the utilization of their analytics workloads and to help meet strict SLAs.

Concurrently, infrastructure cost reduction was becoming a priority: increased data ingest rates and compute requirements were resulting in a 100 percent year-over-year growth rate—causing the company's server capacity to double annually. The I&O team understood that they were on an unsustainable path and sought out Pepperdata Capacity Optimizer.

***“Pepperdata Capacity Optimizer helps optimize the efficiency of our infrastructure to reduce our spending on servers; we also now know what’s needed to achieve optimal performance for each workload.”***

— Vice President, Big Data Solutions Engineering, Global Financial Services Customer



## How Pepperdata Helped the Company Save 30%

The selection of Pepperdata for real-time, automated resource optimization was a natural choice because of its ability to track more than 400 application metrics every three seconds. By automatically identifying underutilized servers and redirecting resources to those workloads in real time, Pepperdata Capacity Optimizer **reduced the company's upward spend on server infrastructure by approximately 30 percent.**

Without needing to make application code changes or apply recommendations, Capacity Optimizer automatically directed the system scheduler to allocate tasks to nodes where more work could be done—increasing the infrastructure's utilization of resources and performance. For this customer, the 30 percent improvement in server utilization and capacity enabled the company to reclaim and repurpose thousands of nodes, thus extracting maximum value from its hardware investment.

Pepperdata Capacity Optimizer also provided a unified solution for resource optimization that eliminated the need for manual tuning, helping the team focus on revenue-generating innovation. With Pepperdata, the I&O team's focus is now laser sharp on managing their growing infrastructure and creating a more efficient Hadoop/Spark cluster environment.

The company considers Pepperdata to be an indispensable and integral part of the company's large-scale data analytics platform as well as its long-term growth and management strategy.

### 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 20,000+ 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)

