



Autodesk Improves Resource Utilization and Reduces Amazon EMR Costs by 50% with Pepperdata



About The Client

Autodesk is a global leader in design and manufacturing software with primary markets in engineering, architecture, construction, manufacturing, media, and entertainment industries.

Challenge:

Autodesk found that scaling Amazon EMR resources to handle its data workloads resulted in runaway costs. Its goal was to reduce costs by 50 percent to accommodate for the continual increase in the company's data processing needs.

Solution:

Autodesk 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:

Autodesk significantly increased resource utilization for its Amazon EMR workloads, reduced Amazon EC2 costs by over 50 percent, and freed its platform data team from manual tuning to focus on revenue-generating innovation.

The Situation: Apache Spark on Amazon EMR Tuning Complications

Autodesk uses Apache Spark on Amazon EMR to process and analyze large sets of big data and turn them into insights. While this approach originally proved effective, workload performance for the company became a significant issue due to resource overprovisioning.

“Spark is notoriously hard to tune correctly. People don’t have time to go into every job. As a result, our entire platform just wasn’t as efficient as it could have been.”

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

Autodesk incorporated optimization methods like instance rightsizing and autoscaling to improve the price/performance of its workloads, but the data team could not keep up with the constant manual config tuning required for optimization.

As the company accumulated more and more data, the increased compute consumption continued to quickly eat through the budget. Each Amazon EMR cluster was consuming two or three times the planned capacity. In 2020, Autodesk saw its data processing needs increase 10x over the previous year.

The company was concerned that if this trend of doubling capacity and overprovisioning resources continued, it would be overwhelmed by runaway costs, low latencies, and increased downtime.

Resolution: Real-Time, Automated Resource Optimization for Amazon EMR Workloads

Autodesk turned to Pepperdata Capacity Optimizer real-time, automated resource optimization to automatically increase the utilization of CPU and memory for the company's data workloads, minimize overprovisioning waste, and reduce costs.

The company set a goal to reduce application costs on Amazon EMR by 50 percent—and with the immediate utilization improvement, Autodesk reduced overprovisioning to help it achieve its cost reduction target.



“After adopting Pepperdata, we were able to see more efficient resource utilization. On average, we were saving 50% on our costs because of Pepperdata—because of the decreased waste and more efficient allocation of resources to applications.”

**—Mark Kidwell, Chief Data Architect,
Autodesk Data Platforms and Solutions**

Pepperdata Capacity Optimizer provided the Amazon EMR system scheduler with real-time visibility into actual CPU and memory utilization levels inside Autodesk's data applications to enable more intelligent resource allocation decisions—automatically, continuously, and without the need for application code changes or manual config tuning. As a result, Autodesk realized a 50 percent reduction of Amazon EC2 instance hours.

Pepperdata Capacity Optimizer enhanced the efficiency of the cloud autoscaler by ensuring new instances were only provisioned when existing instances became fully utilized. The company could run more applications without adding additional hardware and personnel to tune them—saving both money and time for its engineers.

Pepperdata Capacity Optimizer's observability dashboards also provided Autodesk with thousands of application-level metrics in one aggregated view so the data team could quickly and accurately diagnose, troubleshoot, and resolve both cluster-wide and low-level application issues without having to cross-reference multiple dashboards.

After implementing Pepperdata, Autodesk significantly increased resource utilization for its Amazon EMR workloads, reduced Amazon EC2 costs by over 50 percent, and freed its platform data team from manual tuning with a real-time, automated solution that pays for itself.

“Pepperdata allowed us to significantly increase capacity for our Amazon EMR workloads and reduce our Amazon EC2 costs by over 50%. We can focus on our business, while they optimize for costs and performance.”

**—Mark Kidwell, Chief Data Architect,
Autodesk Data Platforms and Solutions**

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.

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



Start a Free PoV
www.pepperdata.com



Send an Email
eval@pepperdata.com

