

Operator Benefits

- Maximize your infrastructure investment
- Achieve up to 50 percent increase in throughput performance
- Ensure cluster stability and efficiency
- Avoid overspending on unnecessary hardware
- Reduce time spent on capacity planning
- Run more jobs concurrently on your existing infrastructure

Develop Benefits

- Run more jobs faster
- Access additional cluster capacity
- Spend less time in backlog queues

With Capacity Optimizer, Enterprises Can:

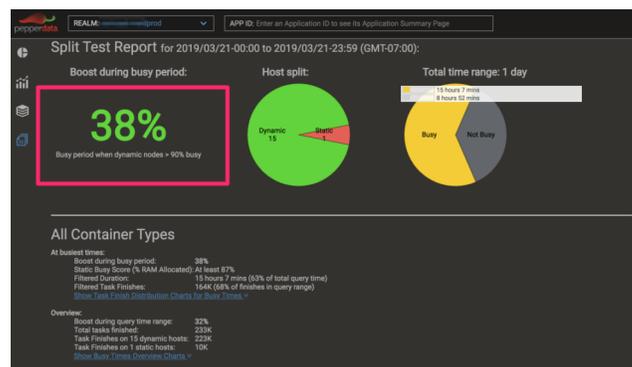
- Maximize your existing infrastructure investment. Increase hardware utilization and eliminate or reduce millions of dollars in hardware costs.
- Auto-tune cluster resources for peak efficiency. Rightsize resource allocation based on actual real-time capacity.
- Eliminate inefficiencies and bottlenecks to ensure infrastructure stability.
- Recapture wasted capacity to optimize cluster resources.
- Run more applications.
- Get the most out of your infrastructure investment.
- Be confident that all YARN capabilities will remain intact.

Continuously Improve Capacity Utilization of Clusters Without Manual Tuning or Intervention

As analytics platforms grow in scale and complexity on-prem and in the cloud, maintaining efficiency is a critical challenge, and money is being wasted. Pepperdata® Capacity Optimizer automatically tunes and optimizes your cluster resources, recapturing wasted capacity so you can run more applications and get the most out of your infrastructure investment.

Improve Cluster Throughput up To 50%

By monitoring the entire infrastructure in real time, including hardware and applications, and leveraging ML with active resource management, Pepperdata Capacity Optimizer identifies where more work can be done and adds tasks to servers with available resources. On a typical cluster, Capacity Optimizer makes thousands of decisions per second, increasing typical enterprise cluster throughput by up to 50 percent. Even the most experienced operator dedicated to resource management can't make manual configuration changes with the precision and speed of Capacity Optimizer.



Doesn't Yarn Scheduler Manage Resources?

YARN ("Yet Another Resource Negotiator") scheduler leverages the resource management capabilities of MapReduce, coordinating consumption and usage reservations to make allocations. Limited by its conservative assumptions about memory usage, YARN under-provisions resources. In addition, YARN does not monitor containers once they start running or adjust in real time based on actual usage.

Pepperdata Capacity Optimizer overcomes these limitations by monitoring actual per-task hardware usage and dynamically making adjustments at the process level, eliminating inefficiencies and bottlenecks, and maximizing resource usage.