Pepperdata 2020 Big Data Performance Report

Enormous Waste—and Enormous Optimization Potential

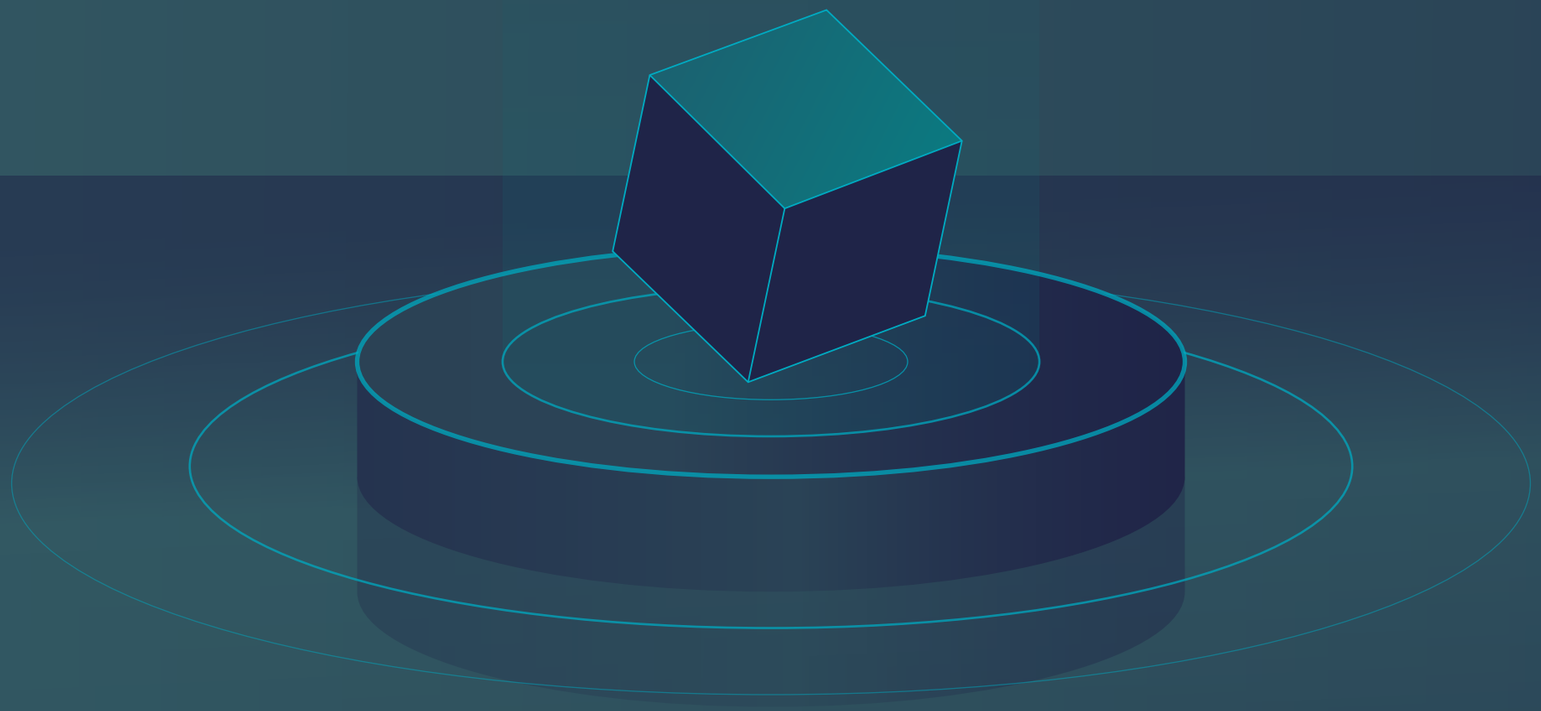

pepperdata

Executive Summary:

IT operations professionals operate in a challenging environment. They are expected to leverage highly complex collections of organizational big data in a way that simultaneously minimizes costs and maximizes performance. Looking to the future of big data, this job is only going to get more challenging.

To shed light on how IT operations teams are dealing with this challenge, Pepperdata has carried out a period of customer research. This report revealed a wealth of insights regarding the condition of enterprise data applications that lack the benefits of observability and continuous tuning. Combined with cloud computing statistics and a more general understanding of big data industry trends, there is much to learn here about the present and future of the data analytics industry.

This report draws on:

- Data from customer clusters—immediately prior to their implementation of cloud optimization with Pepperdata. These clusters represented nearly 400 petabytes of data, on 5000 nodes, running 4.5 million applications.

- Data analysis by Pepperdata field engineers and data scientists.

Our report reveals that, within enterprise applications that are not optimized by solutions that allow for observability and continuous tuning, there exists enormous waste—and enormous potential to optimize applications and cut that waste.

Part One: Wastage

Cloud computing is rapidly expanding. As Statista reports, "in 2020, the public cloud services market is expected to reach around 266.4 billion U.S. dollars in size, and by 2022 market revenue is forecast to exceed 350 billion U.S. dollars."

However, as the cloud expands, so does cloud wastage. As more complex big data applications migrate, the likelihood of resource misallocation rises. This is why, as Gartner reports, "through 2024, nearly all legacy applications migrated to public cloud infrastructure as a service (IaaS) will require optimization to become more cost-effective." Without this optimization, there is guaranteed to be overspend.

In 2019 alone, losses attributed to cloud waste amounted to around $14 billion. One analyst estimates that "$17.6 billion in cloud spend... will be completely wasted this year [2020]." Another analyst estimates that the cost of idle resources alone stands at "$8.8 billion every year." Another major component of wastage is overprovisioned infrastructure; around 40% of all instances are oversized.

When it comes to wastage, failures are important. Job failures can cause serious performance degradation, and consume significant computational resources.

In our unoptimized dataset, we see a wide range of failure rates across clusters. Some clusters will fail above 10%, while others are close to 0%. Spark applications tend to fail more often than other applications.

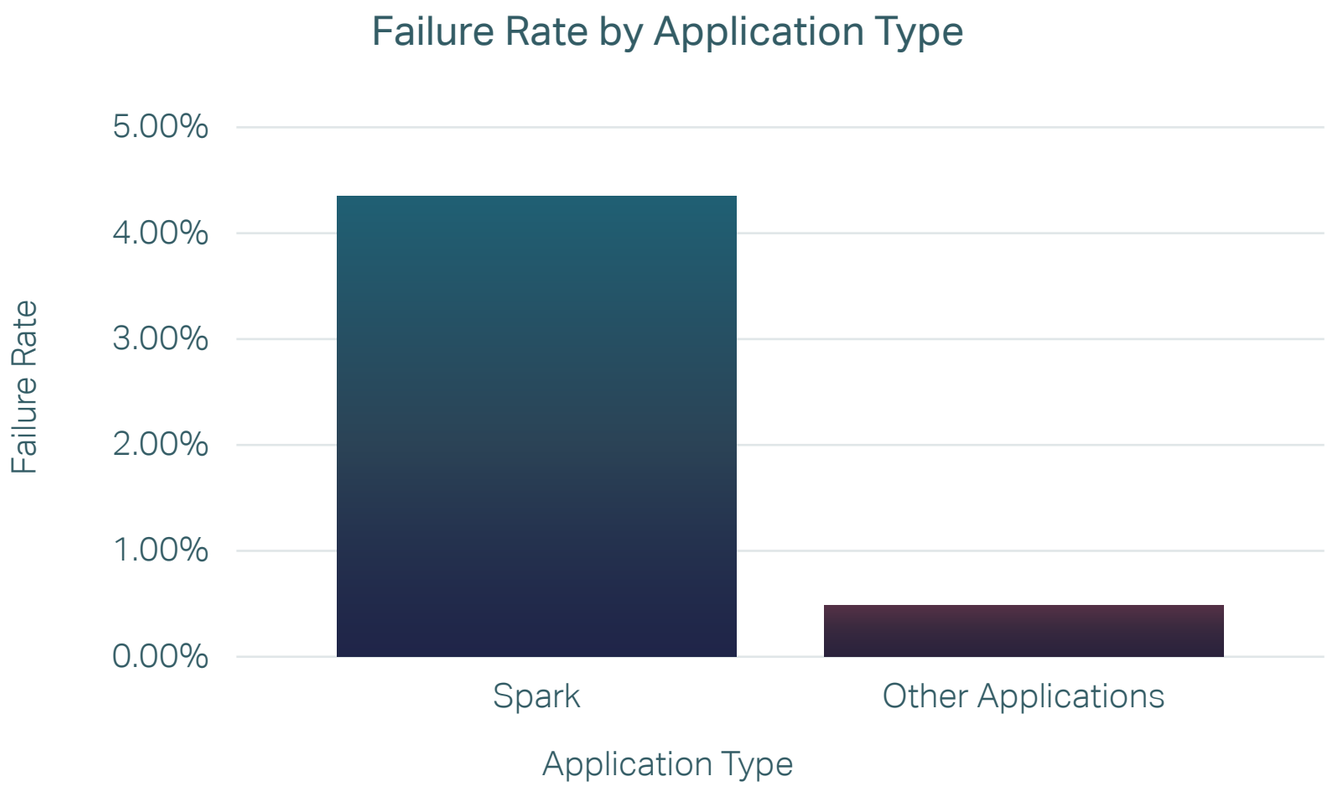

Prior to implementing Spark optimization: Across clusters, within a typical week, the median rate of maximum memory utilization is a mere 42.3%. The underutilization here represents two states: not enough jobs running to fully utilize the cluster resources, or the jobs are wasting resources.

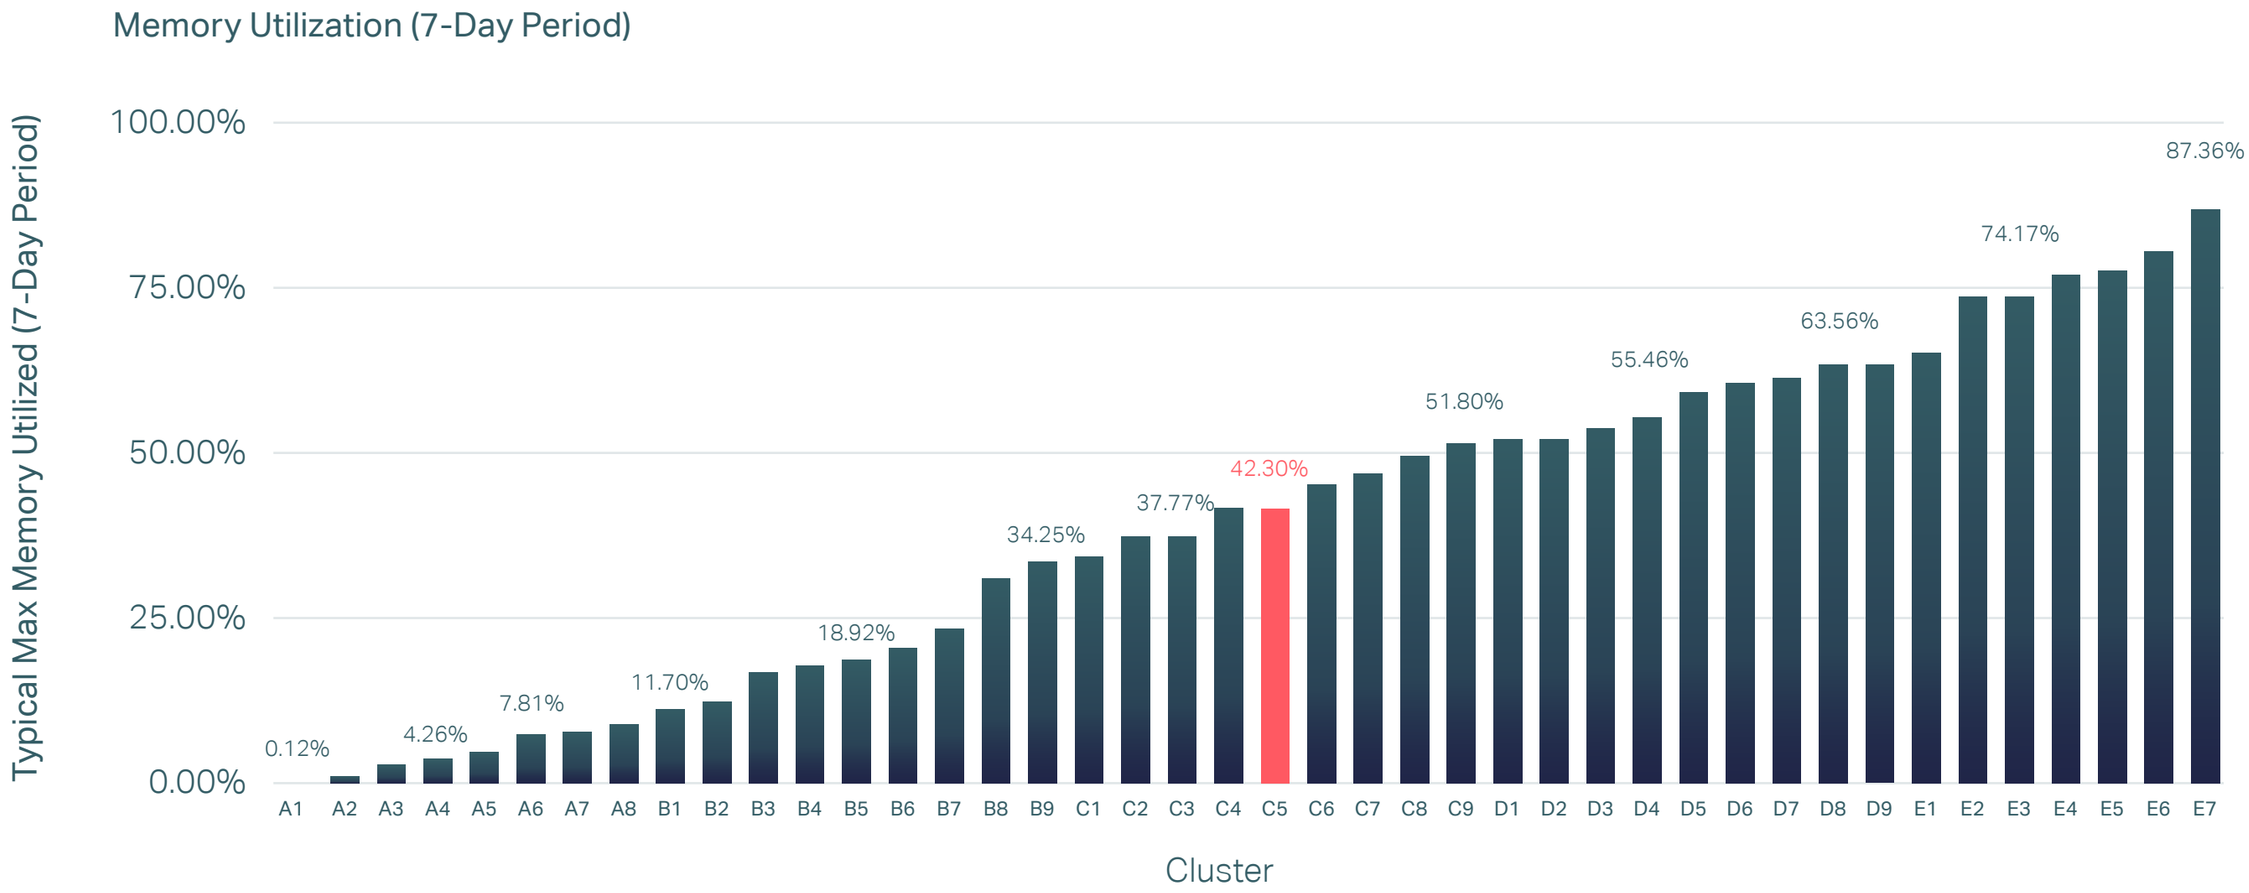

Prior to implementing cloud optimization, when comparing jobs used and jobs wasted, the average wastage across 40 large clusters is 60+%. In the highest cases, it is over 75%.

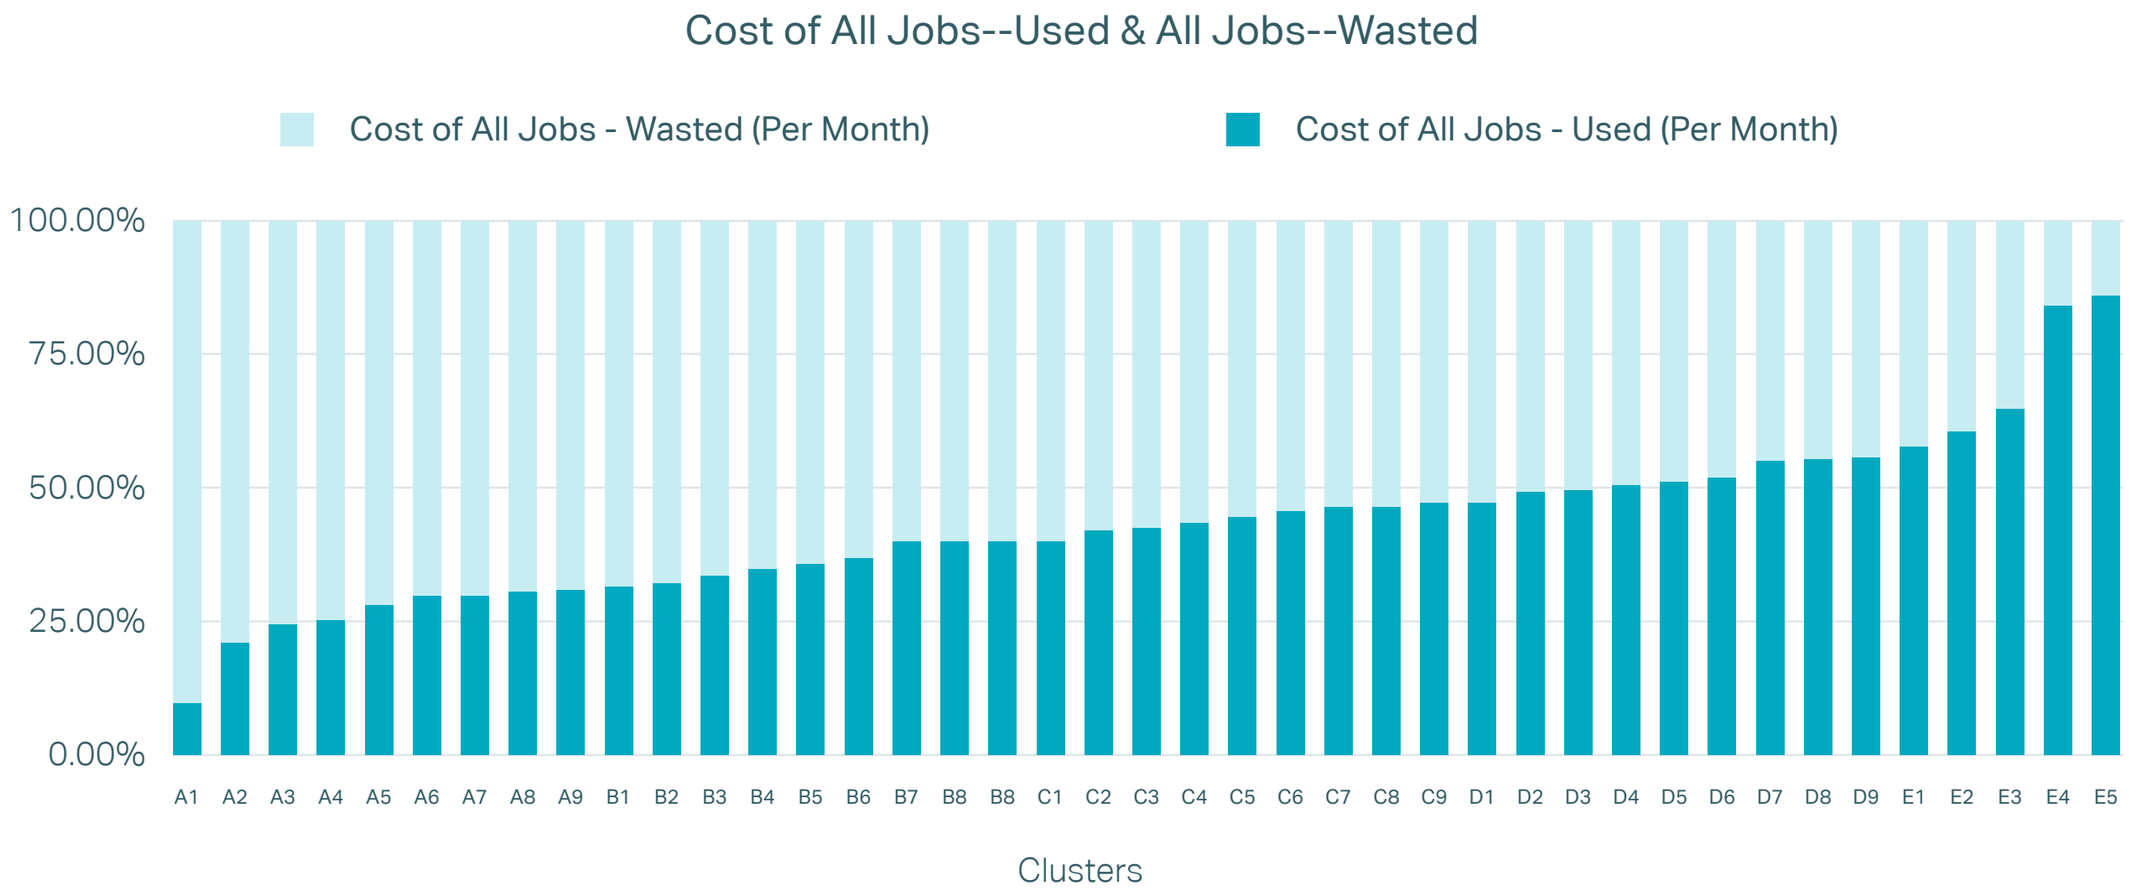

This wastage takes an interesting form. Typically, with 95% of jobs, there is little wastage. Major wastage is often in only 5% to 10% of total jobs. This is why optimization is inherently such a needle-in-a-haystack challenge, and why machine learning can be such a help.

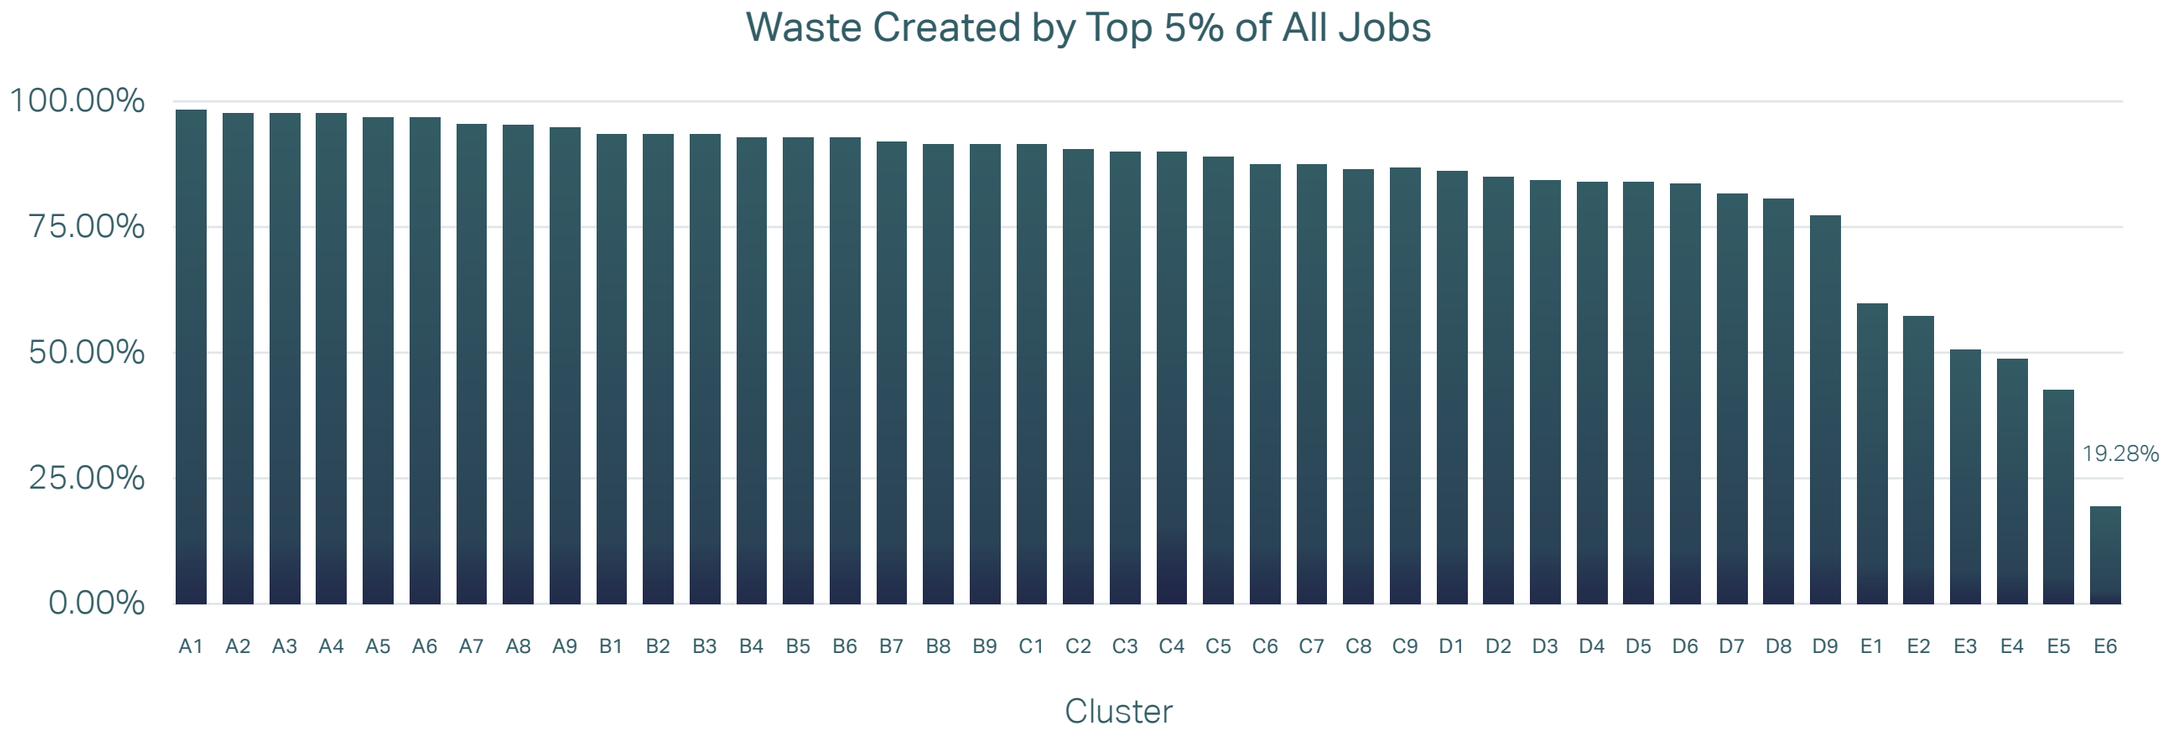

Studies show that ML-powered statistical models can predict task failures with a precision up to 97.4%, and a recall up to 96.2%. Applied to Hadoop, the percentage of failed jobs can be reduced by up to 45%—with an overhead of less than five minutes.

Part Two: Optimization Potential

Cloud optimization can bring big savings. According to Google, even low effort cloud optimization efforts can net a business as much as 10% savings per service within two weeks. Cloud services that are fully optimized and running on extended periods (over six weeks) can save more than 20%.

However, cloud optimization is difficult. Every application has resource requirements that are constantly changing. Big data solutions have to collect and analyze information from all application and associated infrastructure to generate the insights needed to discover optimization potential.

It's critical to dynamically meet these changes in requirements to fully reach optimization potential. Resources have to be deployed to handle the peak of a given application's needs, wherever that peak occurs on its usage curve. If the infrastructure cannot manage that peak, the application will fail.

Big data resource schedulers don't support the frequency of changes required to ramp up to these peaks and then ramp down again. Autoscaling makes large scale changes, when the actual waste is the sum of many smaller instances of the mismatch between allocation and usage.

True cloud optimization must be rooted in machine learning. With the right solutions, applications' resource requirements can be met continuously and automatically, with a speed and dynamism that is impossible using manual methods.

Using a cutting-edge cloud optimization solution, in a typical week, a median user will examine over 54 different unique metrics per cluster. This is the level of detail that is required to maximize optimization efforts.

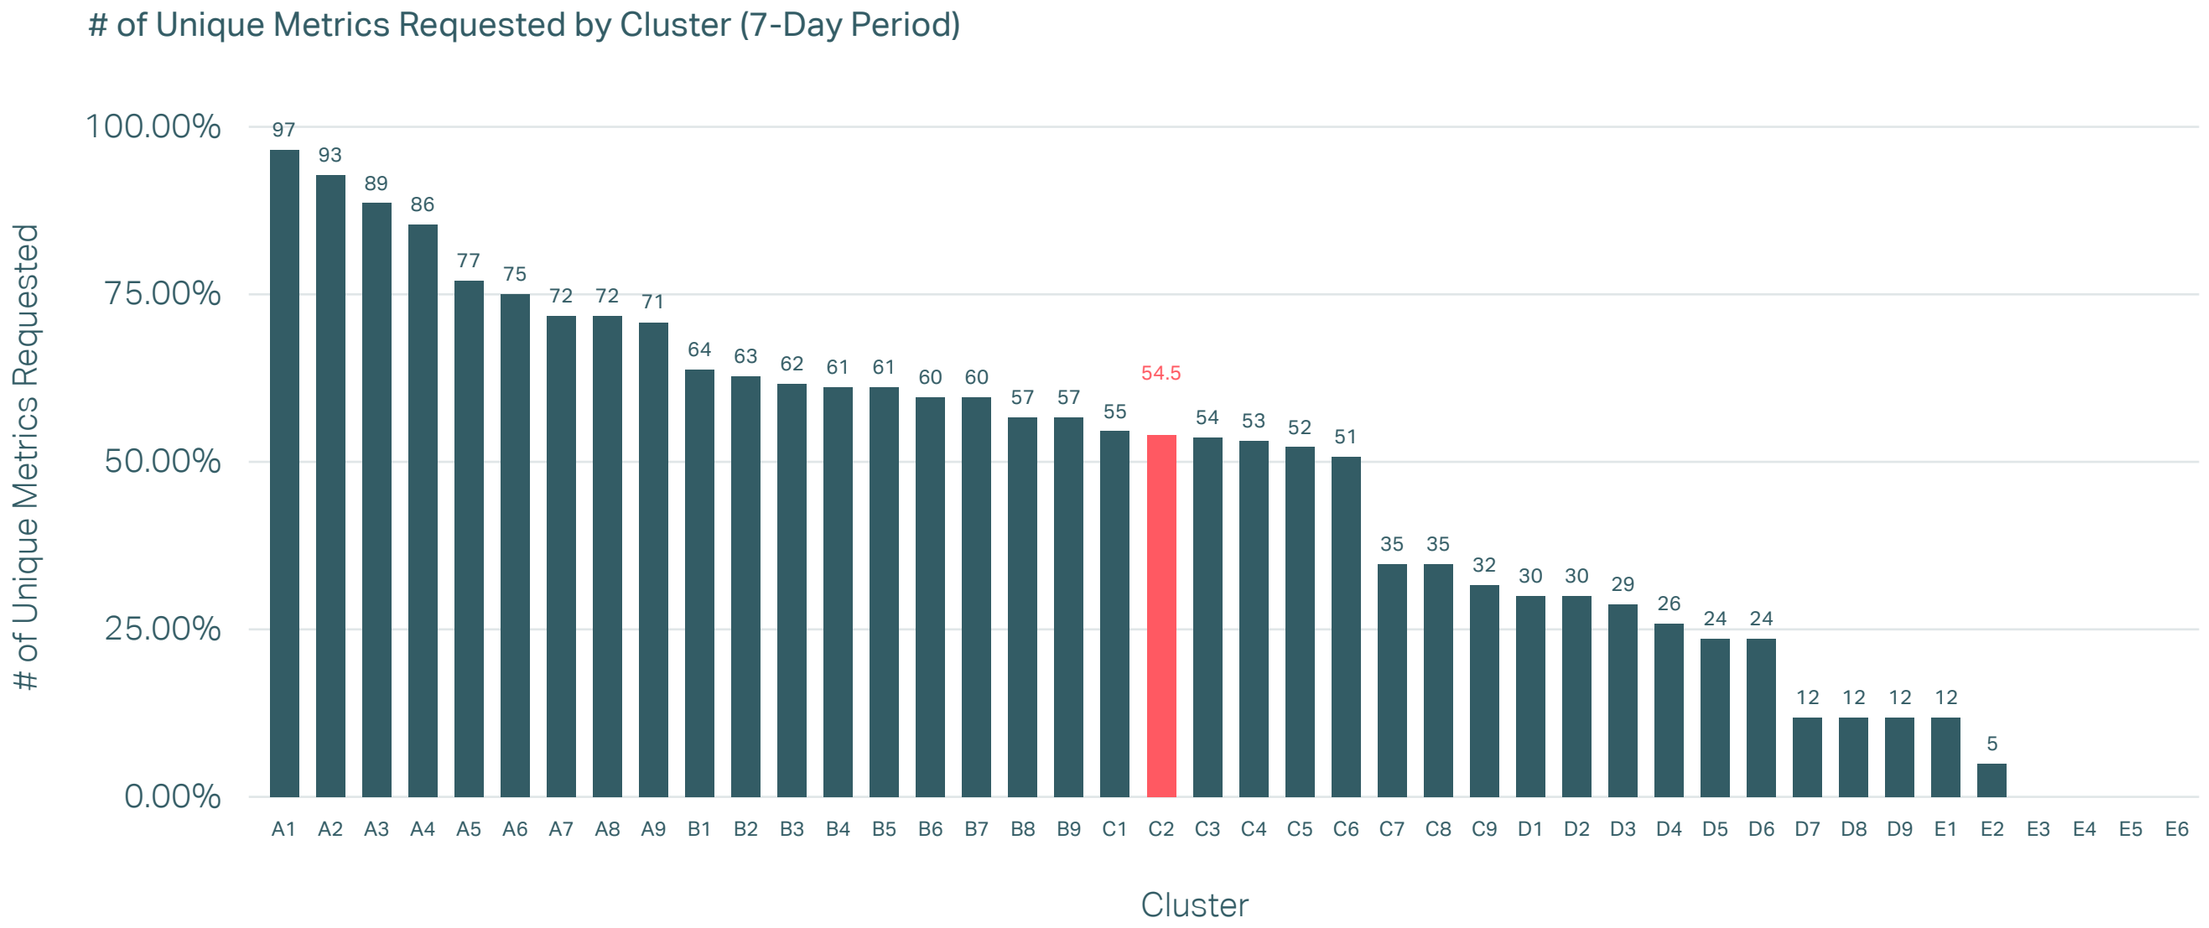

Using the right solution, users can quickly see what is slowing down a particular application, and how different specific factors are impacting bottlenecks.

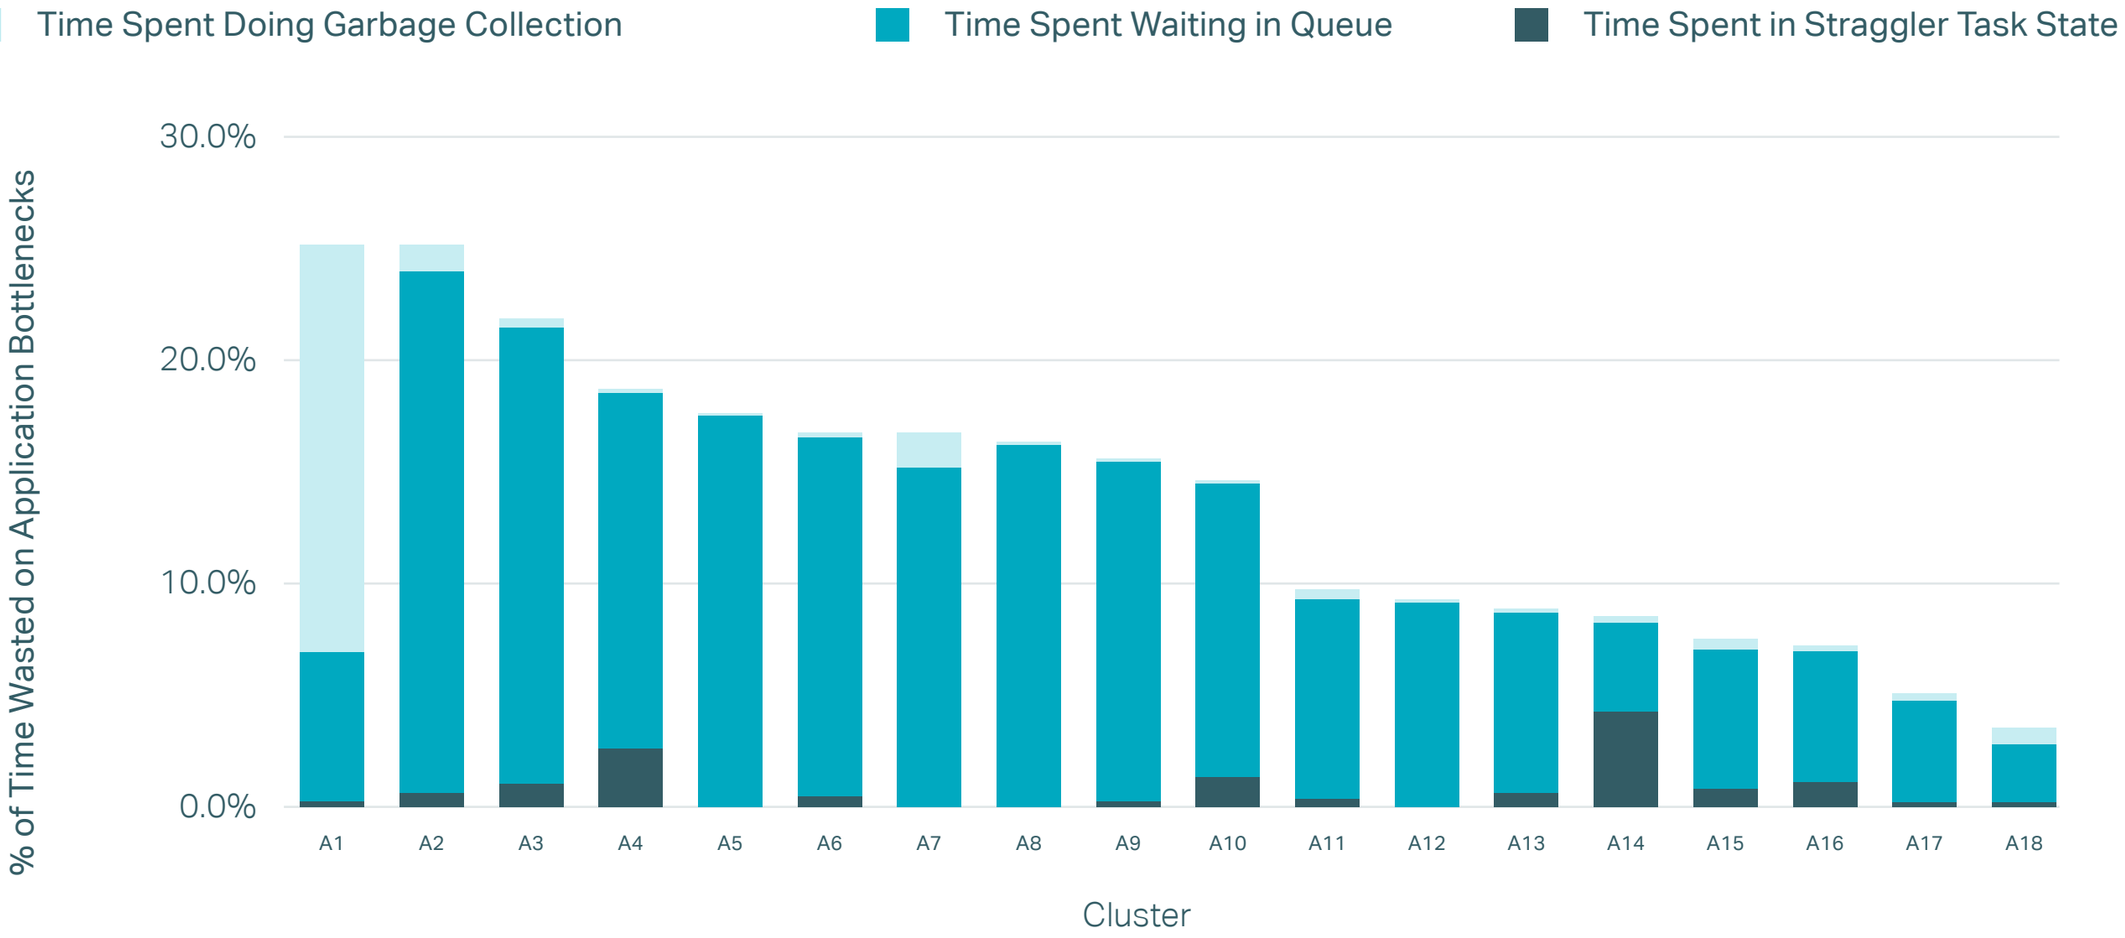

With cloud optimization, three quarters of customer clusters increase throughput and immediately win back task hours. Most enterprises are able to increase task hours by a minimum of 14%. Some enterprises are able to increase task hours by as much as 52%.

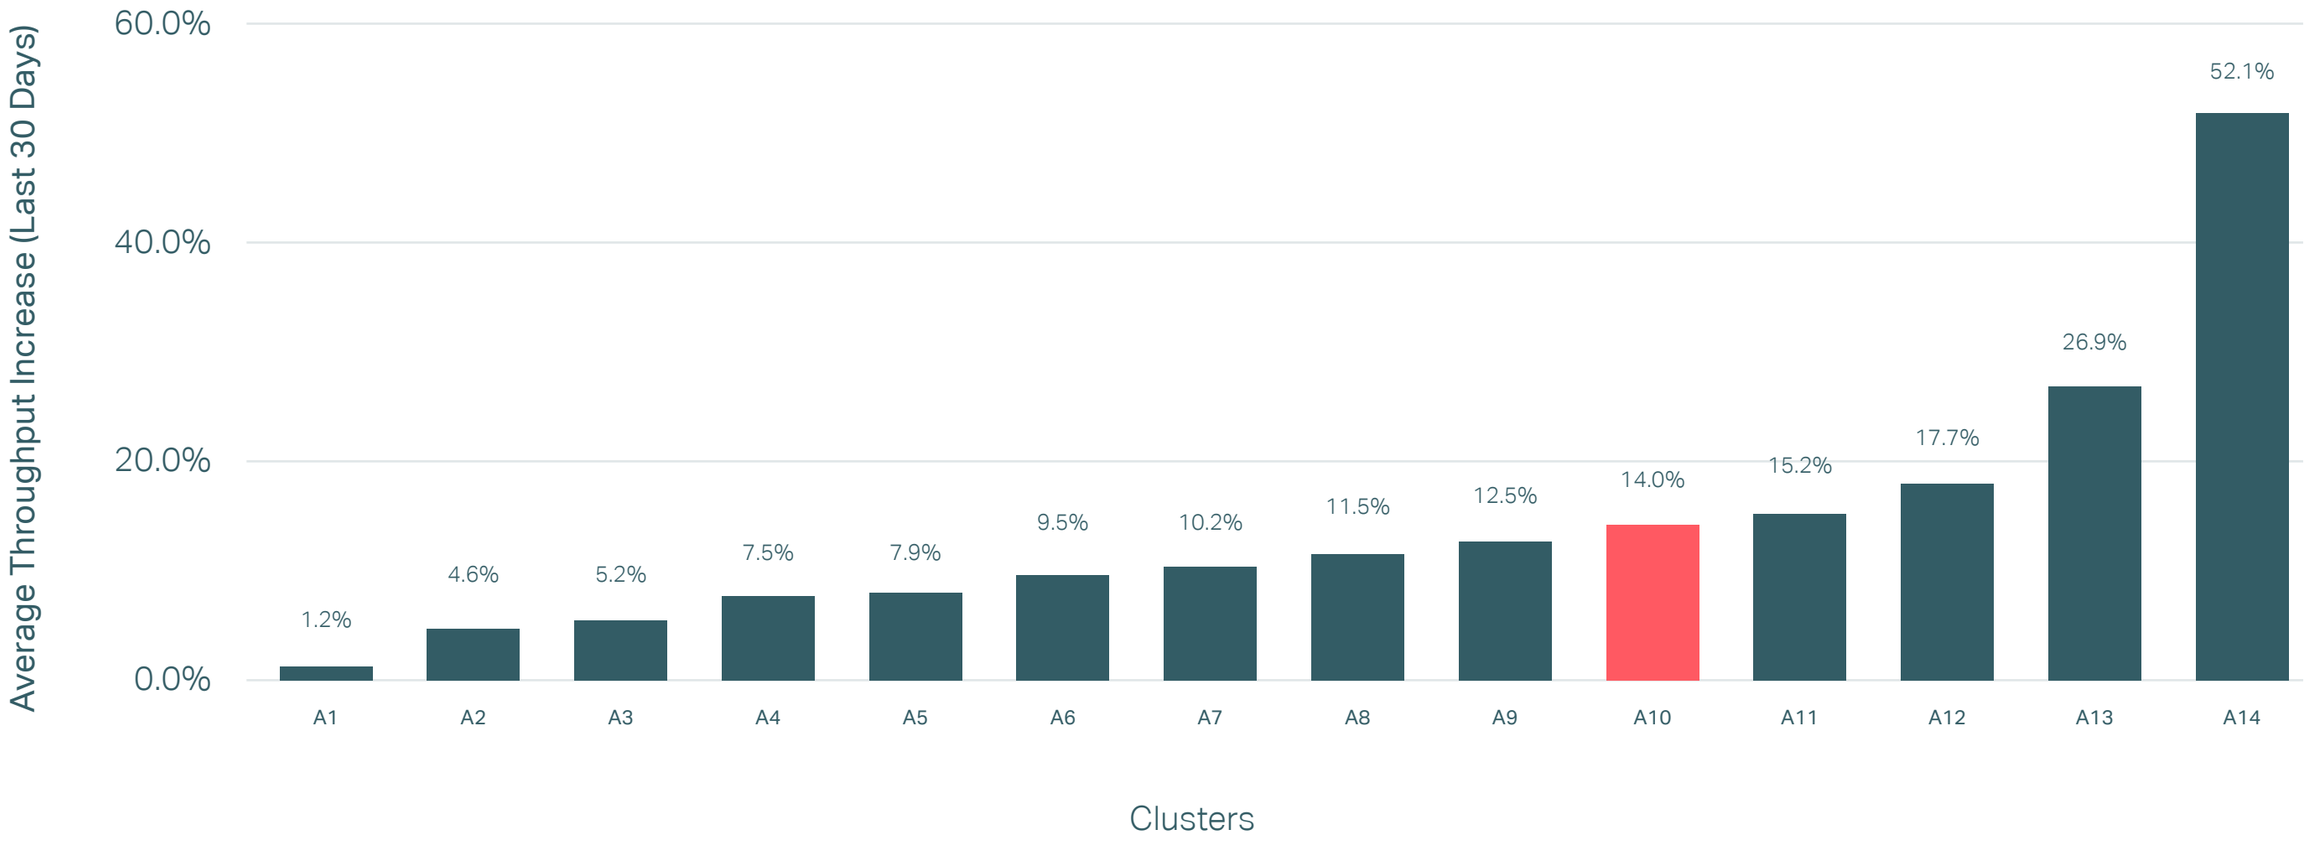

25% of users are able to save a minimum of $400,000 per year.
At the higher end, the most successful users are able to save a projected $7.9 million for the year.

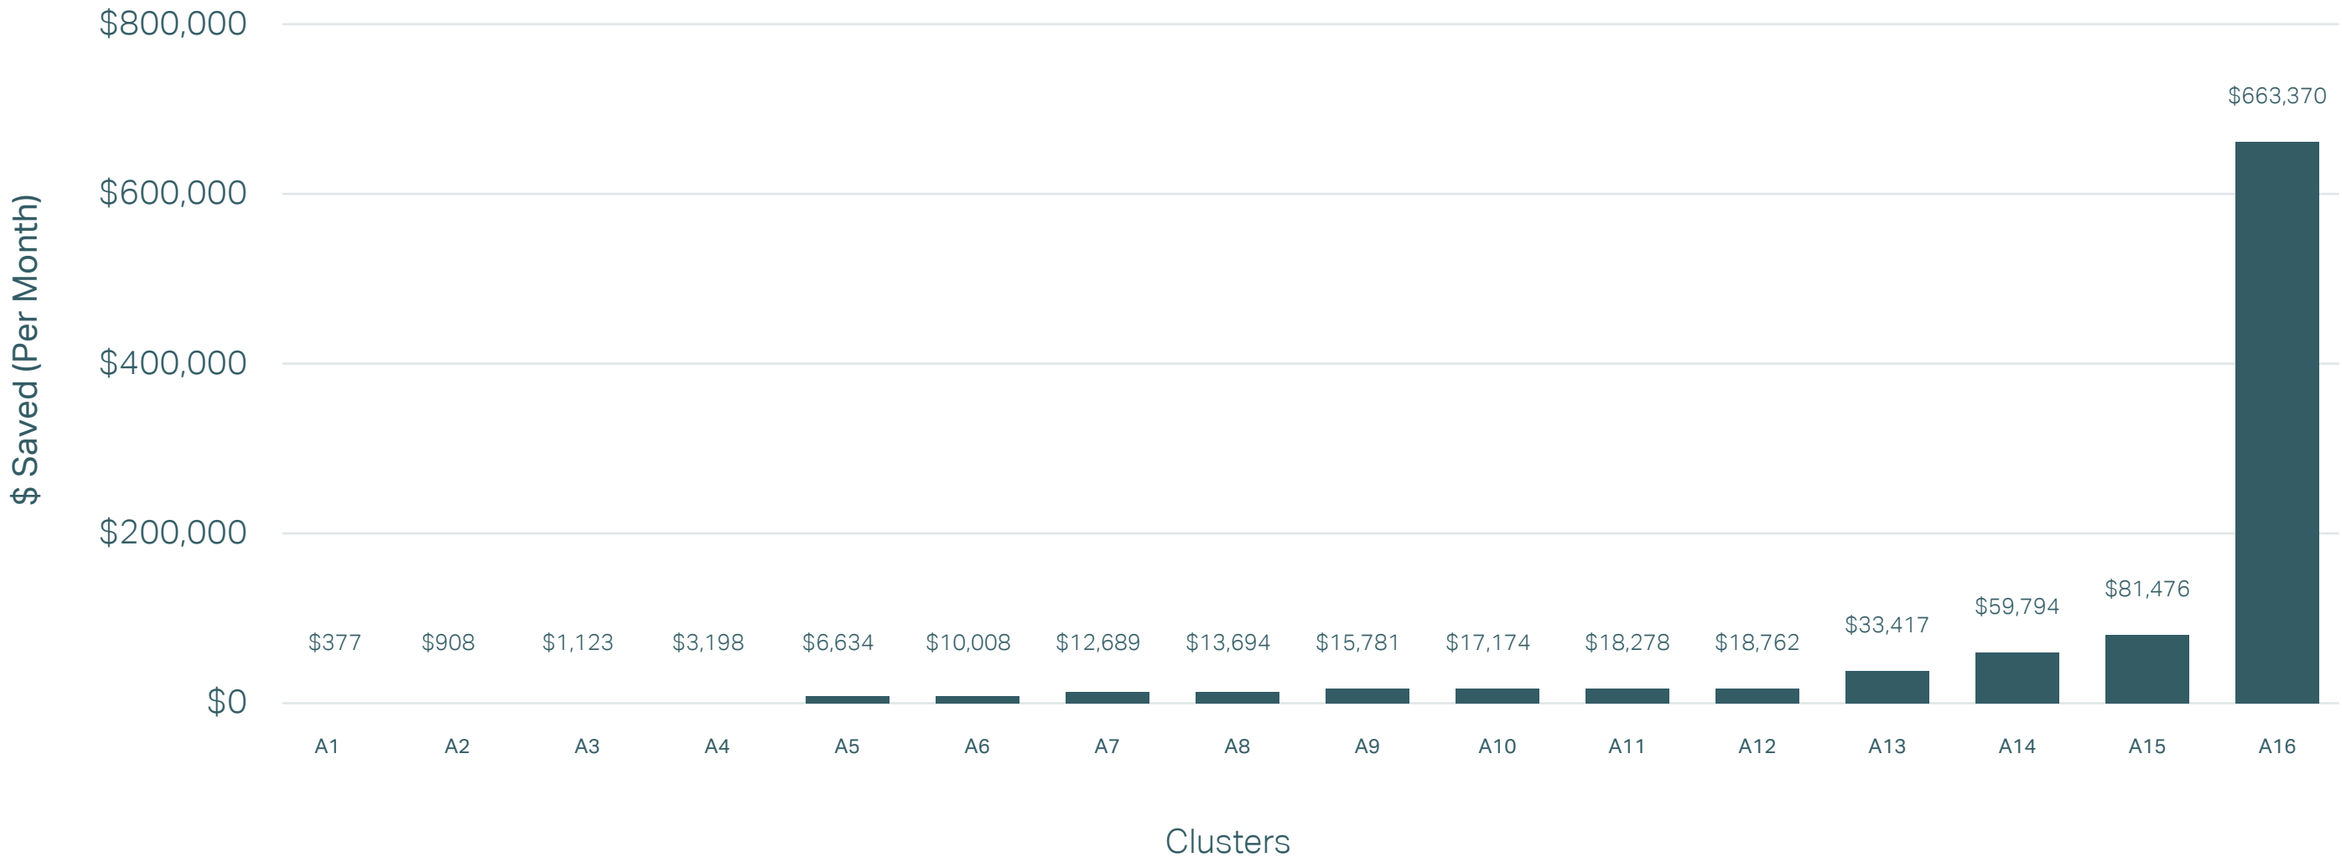

The Future of Applications

As companies continue to migrate to the cloud, operational expenses continue to rise. The cost of unoptimized IT operations processes are set to amount to over $17.6 billion for 2020. Much of that waste is preventable.

Our data shows that big data cloud optimization methods can reduce cost by as much as 20% per service. Many clusters can win back as much as 52% of task hours. Smaller organizations can net $400,000 in annual savings, and larger organizations can save as much as $7.9 million a year.

To cut the waste out of IT operations processes and achieve true cloud optimization, enterprises need visibility and continuous tuning. This requires machine learning and a unified analytics stack performance platform. Such a setup equips IT operations teams with the cloud tools they need to keep their infrastructure running optimally, while minimizing spend.

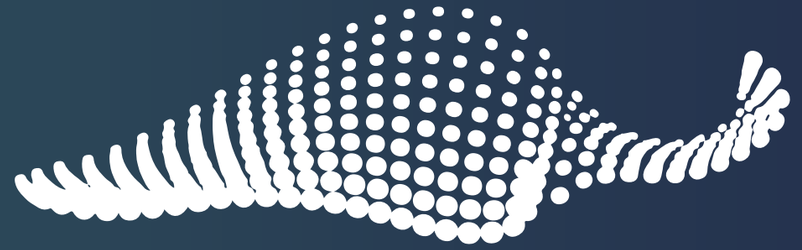

pepperdata®

pepperdata.com