

Cloud Cost Optimization 2.0

Embracing AI, automation and contextual business insights to get more value from cloud computing

Time to retire reactive, inefficient cloud cost optimization strategies and embrace a new approach – combining automation and AI with deeper insights to help organizations achieve superior results.



Introduction

Organizations adopt cloud technology for more convenience, scalability, and cost efficiency. Unfortunately, they rarely get all three. Ambitious teams moving to the cloud and accelerating innovation can foster a virtuous cycle.

However, abundant complexity, unlimited scale, and distributed management of cloud services often leads to less transparency around cloud costs and less control over cloud spending.

COST CONCERNS


50% of enterprises acknowledge they are overspending on cloud

Overspending is a common problem




of organizations will overshoot their IaaS budgets, due to lack of cost optimization approaches, through 2020

Cloud optimization is high on enterprises' to-do lists




of organizations pledged to optimize their cloud usage to save money in 2019, up from 58% in 2018


Implications of poor cloud financial management



Public cloud spending continues to grow



How would you summarize the state of your cloud costs?



- Our cloud costs are about where they should be or lower
- Our cloud costs are a little higher than they should be
- Our cloud costs are way too high

How can Cloud Cost Optimization (CCO) 2.0 help teams maintain the elasticity and rapid innovation that makes cloud technology amazing, without breaking the bank?

FinOps is now a necessity

Cloud financial management, as a practice, moved mainstream due to the accelerated adoption, increased complexity and distributed use of cloud services. This has led to the rise of FinOps. According to the [FinOps Foundation](#), “FinOps is an evolving cloud financial management discipline and cultural practice that enables organizations to get maximum business value by helping engineering, finance, technology and business teams to collaborate on data-driven spending decisions.”

Cloud adoption is accelerating. Along with it, more organizations are embedding FinOps into their business culture, positioning themselves to take full advantage of the cloud. FinOps helps organizations address the unpredictable nature of cloud spending, as well as manage and understand cloud costs.



Companies can do more to manage cloud costs

As a company's use of cloud evolves and becomes more complex, the need for a robust and continual cloud cost optimization strategy becomes essential. In general, the higher your cloud spend, the greater your need is for ongoing management and optimization of cloud costs.

Over 75% of the companies evaluated in the [State of FinOps](#) report have a cloud spend of over \$1 million – the threshold where FinOps proves essential. Yet, among the 41% of these teams spending over \$10 million and 16% spending over \$50 million, only 15% of organizations reported that they consider their FinOps program mature or well-optimized. Enter CCO 2.0.

Cloud Cost Optimization 1.0



In the cloud, engineering decisions are buying decisions. Rather than relying on a CTO or CIO to procure new technology, engineers can decide what they need, click and begin using cloud services within minutes. Engineering choices not only impact the architecture and performance of your applications, but also your cloud spend.

When beginning to build in the cloud, engineers may struggle to fully understand the cost of cloud services. Public cloud providers, like Amazon Web Services (AWS) offer free tiers for most services, so initially cost isn't a concern at all. The AWS bill may remain innocuous for months before incurring a charge. This alone leads to many misinformed engineering choices with expensive long-term impacts. Without visibility into actual cost and lacking business context (i.e., how much is too much and what does good look like?) engineers will decide based on what's best for them and/or their applications, not necessarily what's best for your business or bottom line.

Fast forward to your next phase of growth. Cloud usage is no longer within the free tier. Increased cloud spending invariably invites more scrutiny from finance. Beginning to realize the actual cost of their decisions, engineering leaders find themselves in a state of cloud spend shock, or bill panic, and ask, "How can we reduce cloud costs without sacrificing performance and innovation?" Hindsight may be 20/20, but reactive thinking is the achilles heel of CCO 1.0.

The vicious cycle of cloud costs

CCO 1.0 emerged due to the need for more understanding and control over cloud costs. Organizations using brand-new cloud technologies needed to figure out exactly how much they cost—now and in the future. Then, they needed to determine the best way to go back, reduce and optimize those costs. As noted, CCO 1.0 is reactive. (Think: optimization reports and never-ending checklists of recommendations.)

At a high level, here’s what it looks like, illustrated in figure 1.0:

- Company adopts cloud
- Innovation accelerates (along with cloud costs)
- Cloud costs balloon to the point of causing bill panic
- Engineering focus shifts from innovating to reducing cloud costs
- Once cloud costs are under control, innovation resumes

It’s a brutal, frustrating, and all-too-familiar cycle for many organizations in the cloud. Organizations managing and optimizing cloud costs using 1.0 tactics fall into the same pattern across enterprises, independent of company size, cloud spend, FinOps maturity, and industry.




Figure 1.0 - Vicious Cycle of Cloud Cost



The downsides of CCO 1.0

This cyclical nature of cloud costs forced organizations to adopt best practices for cutting cloud spending and reducing waste. As a result, CCO 1.0 tools and tactics emerged.

CCO 1.0 tools promise to “forecast” cloud usage, estimate savings potential, and recommend optimization opportunities. In reality, these tools provide inaccurate forecasting, overestimate savings, and push an endless stream of optimization tasks on your engineering teams.

The outcome? Rather than building and innovating, teams revise forecasts and manually address optimization recommendations, neither fully achieving optimal engineering performance nor cloud cost efficiency.

Where CCO 1.0 falls short

Despite best intentions, CCO 1.0 practices fail. Organizations perpetually stuck in the vicious cycle of cloud costs experience:

Resource strain

Engineering resources are no longer engineering. Teams are repurposed to tag, track cloudcosts/usage, manually manage reservations, analyze and allocate cloud costs, and more.

Culture misalignment

Engineering teams prefer democratization of the cloud. Finance prefers centralized IT purchasing and spending guardrails. Without the right tools in place to support a centralized and democratized model, teams experience friction.

Reactive, repetitive behaviors

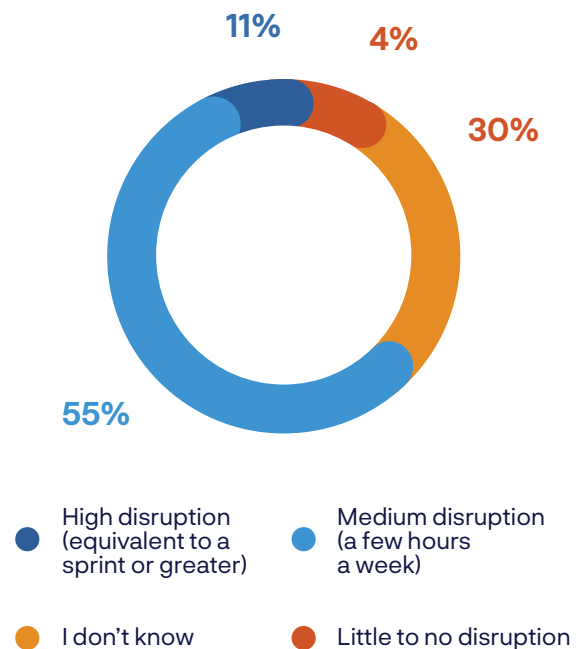
CCO 1.0 organizations prioritize optimizations after cloud cost overruns. Bigger bills are the forcing function for a reduction in cloud spending. Progress is made, but like clockwork, they repeat the cycle.


Incomplete data

Native cloud cost reporting tools, like AWS Cost Explorer, are helpful for figuring your total costs and savings. But, if your entire organization uses the cloud, you need to know unit costs, as well as allocate and understand cloud costs by team, department, application, and user.

Only 3 Out of 10 Organizations Know Exactly Where Their Spend Is Going

In the last 12 months, have cloud costs disrupted your engineering objectives?





Navigating discount instruments in CCO 1.0

Let’s illustrate how a team might leverage discount instruments like RIs and Savings Plans in a CCO 1.0 world.

- Forecasting costs:** CCO 1.0 requires that you’re able to forecast what your cloud spending picture will look like, even if you’re unsure of what your actual usage may be. In short, it’s an educated guess.
- Aligning stakeholders:** In many organizations, this looks like constant coordination between finance and engineering stakeholders to justify and agree on a long-term purchasing plan.
- Ongoing engineering investment:** The engineering team must consistently be involved in the planning, procurement, and management of RIs and Savings Plans.

This process is suboptimal because it can be difficult — if not impossible — to predict the cloud resources your engineering team will need in the future. It’s largely a reactive process that limits agility when done wrong. Not to mention, it’s tough to know how much you’ve saved in the end, or benchmark savings over time.

Many organizations, even those with sophisticated FinOps functions, overestimate savings performance. In reality, CCO 1.0 organizations achieve an Effective Savings Rate (ESR) of less than 20%. For comparison, the top two percent of AWS cost optimizers consistently achieve an ESR of 40% or more.

Effective Savings Rate (ESR)

Is an objective FinOps metric that can be used to measure savings performance of different cloud cost optimization tools, methods, and cloud providers. ESR is expressed as a percentage, representing the discount off the on-demand equivalent (ODE) rates for cloud services.

$$ESR = 1 - \left(\frac{\text{Actual Cloud Spend}}{\text{On-demand Equivalent}} \right)$$

Actual Spend with Discounts is what you paid with RIs and Savings Plans, amortizing any upfront charges.

On-Demand Equivalent (ODE) Spend is what you would have paid if no discounts were applied.

Continue reading about Effective Savings Rate »





Introducing CCO 2.0

ICCO 2.0 is a decisive shift from old, reactive and inefficient cost optimization, to a more intelligent, automated, and continuous approach. Adopting a CCO 2.0 mindset prevents organizations from getting stuck in the cycle of cloud costs. DevOps leaders are empowered to make smarter infrastructure and application architecture decisions based on contextual business insights, not merely top-level cloud usage and cost data.

With CCO 2.0, FinOps teams centrally manage discount instruments, leveraging software that autonomously adapts commitment coverage to usage. Dashboards promote shared visibility into granular cloud costs, giving way to greater financial accountability. Finance, engineering, and business leaders align around shared cloud spending goals. Finally, teams are able to objectively, accurately, and proactively measure the impact of cost optimization activities on cloud costs and savings performance.

How to make the transition to CCO 2.0

Begin by establishing a centralized FinOps function that can champion departmental alignment, define policies, and promote adoption of cloud financial management best practices throughout your organization. Your FinOps program should source stakeholders from finance, engineering, product, and business operations. This centralized FinOps function must enable decentralized decision making—providing DevOps with insights and instructions on how to build intelligently.

CCO 2.0 supports centralized management of cloud costs and savings, while distributing information and providing visibility to developers.

1. Establish finance-aware culture

Everyone in the organization using cloud resources should share responsibility for costs and savings performance. Stakeholders need to make buying decisions based on actual business value.

Implementing a FinOps culture will shorten the feedback loop between purchasing cloud resources (by engineering) and the financial impact, enabling engineers to make cost-informed decisions. Cloud financial data should be visible, shared throughout the organization, and reviewed regularly.


Finance-aware engineers are empowered to evaluate and make decisions based on speed, performance, and long-term business value.

“Decentralize using less (avoiding costs), centralize paying less (reducing rates).”

- J.R. Stormont, executive director of the FinOps Foundation



Correlation between engineering ownership and cost outcomes



2. Automate manual optimization tasks

Rather than discrete, point-in-time optimizations, CCO 2.0 is ongoing. Manual cost optimizations, like managing cloud discount instruments, should happen autonomously, guided by real-time telemetry from your cloud computing environment.

To overcome human limitations and achieve optimal results, DevOps and FinOps leaders should leverage machine learning and AI to automate and eliminate manual tasks.

3. Focus on unit metrics and business value

CCO 2.0 helps you unlock the business value of your cloud. Beyond top-level cloud cost metrics, drill into unit metrics—cost per customer, product, feature, transaction—measurements meaningful to your business and bottom line.

Cloud cost reporting should provide intelligence, rich insights, and granular visibility into the business impact of your organization’s cloud spend and optimization activities.

By focusing on unit cost instead of total cloud cost, teams can see the forest through trees (fig 2.0). For example, your team may discover certain features are underpriced (citing unit costs, specific to that feature) and propose a pricing structure more beneficial to your business. These sorts of changes aren’t about cost-cutting alone. They can also drive top-line growth.




Fig 2.0 - Total Cloud Spend vs. Unit Cost of Cloud

Key Takeaways



Organizations struggle to break the cycle of cloud costs, which includes cyclical acceleration, bill panic, and stalled innovation. This cycle, caused by a failure to implement ongoing cloud financial management processes, impedes growth.



Engineering and finance lack alignment when it comes to cloud spending goals, which illustrates the need for a culture shift. When it comes to controlling cloud costs, engineering and finance should work together, maintaining focus on speed of innovation and optimal cost efficiency.



Traditional cloud cost reporting and management tools primarily display data in arrears. CCO 2.0 tools enrich data with insights, enabling teams to make informed, forward-looking decisions based on contextual cost metrics, like unit cost.



Manual cloud cost optimization tools, such as software that provides reporting and recommendations, are limited by human ability. Achieving optimal outcomes requires automation and AI. CCO 2.0 technology unlocks human potential and eliminates cost inefficiencies in ways people simply cannot.



Effective Savings Rate (ESR) is an objective FinOps KPI that can be used to benchmark, track, and understand the impact of cloud cost optimizations. ESR measures the success (ROI) of your cost optimization activities.

Get started with CCO 2.0



ProsperOps and CloudZero provide organizations the tools, technology, and guidance needed to make the jump to CCO 2.0. With autonomous cost optimization, instant showback, and intelligent reporting tools, our platforms simplify cloud financial management.

CCO 2.0 Next Steps

- Embrace cost optimization tools that leverage automation and machine learning to achieve optimal results.
- Cultivate a finance-aware culture that centralizes rate optimization and democratizes cloud purchasing decisions.
- Establish a FinOps function sourcing stakeholders from finance, engineering, product, and operations teams.
- Implement cost reporting solutions that provide intelligence and insights, adding context around cloud spend, unit economics, and business value.
- Benchmark your ESR and current savings performance at effectivesavingsrate.com, then track it over time to understand the impact of CCO 2.0 on savings performance.

About **CLOUDZERO**

CloudZero is the cloud cost intelligence platform that puts spend into the context of your business. By aligning engineering, infrastructure, and finance teams around metrics like cost per product feature, customer, and development team, CloudZero enables better strategic decisions, improved unit economics, and efficient spending. Trusted by top cloud-driven companies like Rapid7, Ping Identity, and Malwarebytes, CloudZero works with organizations of all sizes to take the next steps toward cloud cost maturity.

About ProsperOps

ProsperOps is a leader in AIOps for cloud financial management. Founded in 2018, ProsperOps provides an intuitive and autonomous cost optimization experience for Amazon Web Services (AWS) that automatically manages AWS discount instruments to maximize compute savings and minimize commitment risk. By removing the effort, latency, and risk associated with manually managing rigid, long-term discount instruments, ProsperOps simplifies cloud financial management. ProsperOps is a founding member of the FinOps Foundation and a FinOps Certified Platform, AWS Advanced Technology & ISV-Accelerate Partner, and 2021 Gartner Cool Vendor in Cloud Computing.