SAVE MORE

RATE OPTIMIZATION:

The Cloud Cost-Saving Strategy You Need to Know

White Paper I Aditya Datta

LOWER CLOUD COST AND BOOST YOUR INFRASTRUCTURE MARGIN.

Schedule a call with our cloud experts.
Try our calculator to see how much you can save.

Rate optimization can reduce spend by up to 30%.

ABSTRACT

Even with manual efforts to reign in cloud costs, most organizations consistently overspend, leaving untold dollars on the table every month. Rate optimization can reduce spend by up to 30% with zero impact to infrastructure configuration – but the vast majority of companies aren't taking full advantage of this key FinOps strategy.

Read this white paper for a better understanding of how rate optimization helps teams avoid overspending, improve infrastructure margins and reduce OpEx. You'll get key insights into savings tools specific to top cloud providers (AWS, Microsoft Azure and GCP), and learn how working with a trusted third-party partner can deliver more savings than in-house efforts alone.

CHALLENGE: SKY-HIGH CLOUD SPEND

Migrate to the cloud, ASAP. It's been the mantra for a decade, and for good reason: by now, no one in IT needs a breakdown of the benefits of cloud versus on-premises computing and data storage.

But when it comes to cost management, one of those long-held "why cloud" pillars, there's a dark side: after a while, operational costs tend to balloon. In fact, Gartner estimates companies overspend on cloud by up to 70% in the absence of consistent or ongoing planning and managment.

Some of the main culprits for cloud costs escalating over time include:

- Lack of understanding of all available savings instruments
- The administrative burden of manually analyzing hundreds of small line items every month
- The constantly-moving target of usage
- Limited accountability among teams

WHAT EXACTLY IS RATE OPTIMIZATION, AND WHY SHOULD WE BE DOING IT?

Enter rate optimization, the practice of ensuring that you're consistently paying the best price for the cloud resources you're using. Rate optimization can lower total cost of ownership (TCO) of cloud services through savings instruments like reservations, savings plans, enterprise discount programs, and private pricing agreements.

RATE OPTIMIZATION IS A MUST-HAVE COST-REDUCTION STRATEGY FOR A COUPLE OF KEY REASONS:

Easy on IT.
Rate optimization can be performed centrally, by individuals or teams, with no involvement from IT, Engineering or DevOps teams.

Significant savings.
Cloud consumption eligible for savings instruments typically makes up the largest portion of the cloud bill, making rate optimization an effective tool.

Zero impact
on infrastructure design or performance. Implementing an effective rate optimization strategy will not impact cloud configuration or performance of applications.

FinOps Context for Rate Optimization

Let's get basic for a minute: cloud costs are a function of how much you use multiplied by the price of that usage type. That's where we get the simple cloud equation:

Usage x Rate = Cost

This means that FinOps teams have just two tools available to reduce cloud cost: usage and rate reduction. From FinOps experience, we know that rate reduction works best when centralized, while usage reduction works best when decentralized. "Centralized" here means that a single person or a small team can potentially perform the activity without relying on IT engineers for support.

TO ILLUSTRATE, TAKE A LOOK AT THIS CLOUD LINE-ITEM EXAMPLE FROM AWS:

	Linux On Demand cost	Linux Reserved cost 1 year	Linux Reserved cost 3 year
C5.2XLarge	$0.34 hourly	$0.24 hourly	$0.162 hourly

THE BOTTOM LINE:

Monthly cost for 30 days without reservation:
$0.34 x 24 hours x 30 days = $244.80

Monthly cost with 3 year for 30 days with reservation:
$0.162 x 24 hours x 30 days = $116.64

Savings of 52%

What Are the Keys to a Successful Rate Optimization Strategy?

Rate optimization can net some very worthwhile savings. There's a learning curve, though, and putting some time into planning ahead can help FinOps teams trim costs more quickly.

Here are three things to consider before you implement rate optimization.

KNOW YOUR KPIs

If you can't measure it, you can't manage it – so be sure to establish goals and key performance indicators (KPIs) before beginning your rate optimization journey.

These are the industry-accepted KPIs you should be using to track rate optimization effectiveness:

- **Coverage:** How much of your on-demand usage is covered by savings instruments?
- **Utilization:** How much of your savings instrument inventory is being applied to on-demand usage?
- **Dollars saved:** How much have you saved over time through rate optimization?
- **Commitment flexibility:** How much can you decrease your commitment at any given time?

UNDERSTAND HOW RATE OPTIMIZATION VARIES AMONG TOP PROVIDERS

Because most companies use a combination of cloud services from AWS, Microsoft Azure, and GCP, it's worthwhile to have a baseline understanding of how rate optimization works with each provider:

Amazon Web Services

Of the major cloud providers, Amazon Web Services (AWS) offers the broadest variety of savings instruments, including Savings Plans, Reserved Instances, Convertible Reserved Instances and Enterprise Discounts. However, there are a few important caveats to keep in mind:

- Not all AWS savings products can be exchanged – so once a purchase is made, the commitment cannot be changed.
- Because some of the Reserved Instances products have a very narrow focus regarding which cloud workloads they apply to, misguided savings purchases are common in the early adoption period, before staff have come up to speed on training and experience.
- The AWS cloud native reservation recommendation tool requires trained staff to compensate by over-predicting future growth.

Amazon Web Services Key takeaways:

- The best savings with AWS are achieved using a combination of Savings Plans and Reserved Instances.
- Effective savings portfolio management with AWS requires training and experience.
- Have systems in place that provide realtime visibility into overall discount and commitment flexibility.

Microsoft Azure

Unlike AWS, Microsoft Azure offers reservations for storage, such as disk and blob storage, in addition to compute and database resources. All Azure reservations can be exchanged for a reservation of the same type, a flexibility that does not broadly exist on the AWS side. In addition, customers can also refund reservations up to $50,000 within a 12-month rolling window.

Advisor, the Azure cloud-native reservation recommendation tool, uses a conservative approach to purchase recommendations. On the plus side, staff new to the cloud can effectively self-manage reservation purchases. However, Advisor will not surface new workloads that can benefit from reservations – a result stemming from the tool's conservative approach.

Microsoft Azure Key takeaways:

- Savings portfolio management is less complex than AWS.
- There's less visibility into which cloud workloads can benefit from reservations, especially for newer workloads.
- Have systems in place that provide realtime visibility into overall discount and commitment flexibility.

Google Cloud Platform

Google Cloud Platform (GCP) offers Committed Use Discounts (CUD) for only five cloud services: Cloud SQL, Cloud Spanner, Google Cloud VMware Engine, Cloud Run, and Google Kubernetes Engine.

Google is the only cloud vendor that offers sustained usage discounts (SUDs) for vCPUs and memory. Although SUD savings are lower than CUDs, a major benefit is that SUDs are applied automatically when cloud workloads run continuously over a longer time period.

Cost Recommendations, the GCP cloud-native reservation recommendation tool, is rudimentary and conservative.

Google Cloud Platform Key takeaways:

- GCP savings portfolio management is the least complex of the major cloud providers.
- Compared to Azure, GCP offers greater visibility into which cloud workloads can benefit from reservations.

UNDERSTAND ONE-TIME AGREEMENTS

In addition to savings instruments and reservations, major cloud providers offer one-time agreements to provide a discount based on overall committed spend over a one or three year period. Here's what you need to know about Enterprise Agreements and Private Pricing Agreements.

Enterprise Agreements
All major cloud providers offer Enterprise Agreements (EA) for high-volume customers. As a best practice, reach out to your account manager to check if or when your cloud usage is eligible for an EA. If you're unfamiliar with EAs, here's a closer look at what they are and how they work.

An EA is a legally binding agreement between two companies that states the discounts the vendor will provide in return for a spend commitment over a specified time period. For example, a customer commits to spending $100M over a 3-year period and receives a 12% discount on all cloud products purchased. If the customer doesn't meet the spend commitment, the remainder typically will be due when the enterprise agreement ends.

Negotiations generally take three to eight weeks depending on complexity. While knowledge around discount amounts related to spend commitment is a closely guarded industry secret, authorities like Gartner provide guidance on what to expect.

Pro tip: Be aware that negotiating skill is key to extracting the most benefit from an EA.

Private Pricing Agreements
A Private Pricing Agreement (PPA) is similar to an EA but is usually limited to a specific cloud service. For example, a customer with blob storage usage of 50 PB over a 3-year period might receive a 16% discount.

PPAs often are amendments to EAs, and an EA can have multiple PPAs. When renegotiating an EA with multiple PPAs, it is best practice to incorporate existing PPAs into the new EA whenever possible.

Pro tip: GCP offers tens of thousands of so-called stock keeping units, or SKU Groups, where each SKU Group can be targeted by a PPA.

RATE OPTIMIZATION IN PRACTICE:

Alkami Sees 15x Savings

COMPANY

Alkami Technology, a leading cloud-based digital banking solutions provider for U.S.-based banks and credit unions.

PROBLEM

After migrating to AWS, Alkami wanted to get the most from their hosting investment while improving internal finance efficiencies.

SOLUTION

Alkami partnered with Cloudwiry to help optimize costs and accelerate the conversion process to cloud native.

OUTCOME

Alkami increased savings 15x in the first month using convertible reserved instances. Cloudwiry's unique Reserved Instances (RI) capabilities enabled Alkami to commit to 95%+ coverage for a 3-year convertible RI plan, exclusively operating within AWS. Another great result? Alkami's FinOps team is now free to focus on long-term cloud economics goals.

> " Cloudwiry's autopilot saves us money and time every day – what used to be a chore now happens daily on autopilot, simplifying cost management and reporting significantly for us. This is truly hard dollars and time saved for our team."
>
> – **Marc Jones**, CTO, Alkami Technology

Is Third-Party Management of Rate Optimization Right for You?

Companies often use a combination of self-managed and third-party-managed rate optimization methods. The capability of in-house teams to manage instruments typically caps out at around 80% of cloud savings.

In contrast, third-party reservation management with a trusted partner can leverage nearly 100% of savings opportunities.

Considering third-party rate optimization management? Keep in mind that the ability of a rate optimization partner to deliver maximum savings depends on the experience, dedication and knowledge of staff, as well as access to industry-leading technology.

Since 2016, Cloudwiry has helped clients accelerate cloud savings through tech-enabled automation, innovation, and deep expertise.

LOWER CLOUD COST AND BOOST YOUR INFRASTRUCTURE MARGIN.

Schedule a call
with our cloud experts.

Try our calculator
to see how much you can save.

View our fireside chat
on rate optimization.

ABOUT CLOUDWIRY

All great companies start with a problem. Ours was maximizing the return on public cloud investments. We started as practitioners in cloud cost management and quickly realized the potential savings are significant. We also realized that implementing and maintaining FinOps best practices is an ongoing effort requiring deep expertise and consistent focus. We saw an immediate need in the market for automated financial management and clear cost insights. Cloudwiry was founded in 2016 to solve that need with a technology-first approach.

cloudwiry.com