

Guide to AWS Billing Data

How to Use Your AWS Billing Data to Optimize Cloud Costs

Contents

PART 1

The Basics 04

PART 2

Consolidated Billing 07

PART 3

AWS Billing Data 09

The Importance of Understanding Your Billing — And Your Cloud

The landscape of business has drastically changed in the last ten years. IT has gone from a line item of support to crucial infrastructure that's essential for businesses to function. At the same time, the rapid evolution of the cloud has moved IT from onsite data centers to cloud-based infrastructures.

As a result, the cloud has become a vital part of organizations, one that reaches out into all aspects of their business. When this happens, the health of your cloud becomes directly tied to the health of your company. The more you understand your cloud and how it relates to your structure and processes, the better decisions you can make about your cloud. When you make better cloud decisions, you increase the efficiency of your cloud, which increases the efficiency of your business.

Using Billing to Understand Cloud Usage

The great promise of Amazon Web Services (AWS) is that you'll only pay for what you use instead of having to invest in data centers that end up underutilized. On the surface, this makes for a fairly simple equation: you use AWS services, then pay your bill for what you use.

But dig a little deeper and you'll find that the cloud "pay for what you use" model also opens the door to increased efficiency and optimization. If you can focus your AWS resources to only be what you need and to take advantage of discounts, you'll end up with a smaller cloud spend.

AWS usage and billing are linked by their very nature. The more you understand your AWS billing, the more you'll be able to optimize your cloud. And the more you take advantage of your billing data to fuel your AWS analysis, the more valuable insights you'll get to make vital business decisions.

This guide is here to help give you a core understanding of AWS billing and to show you how you can take advantage of the wealth of data it offers.

PART 1:

Basics

Billing Basics

AWS billing data can be accessed several different ways. Each has its own pros and cons, but in the end it really comes down to how deeply you want to dig into the data. The spectrum goes from simply viewing and paying your bill all the way to downloading granular usage data that can be used to fuel machine learning recommendations and forecasts.

AWS Billing Console

The AWS Billing Console gives you access to your cost and usage information. There are a variety of AWS tools and products integrated into the Billing Console to help you make sense of your billing data:

- Billing Dashboard gives you a basic overview of your AWS billing and usage, such as your monthly costs to date and your biggest spending services.
- Bills Page has all your past bills for accounting or reference.
- Budgets provides you the ability to set custom budget alerts.

- Cost Explorer lets you take a fuller look at costs than Billing Dashboard does.
- Reserved Instance Reporting shows you core data to track Reserved Instance usage.
- Cost and Usage Reports contain the most comprehensive data about your AWS account.

The various out-of-box tools can be very useful for getting a basic idea about your AWS bill and how it relates to your AWS usage. They help translate your bill from a spreadsheet of data into a graphical representation that makes it easier to see your usage. For example, AWS Cost Explorer includes reports like Monthly Costs by AWS Service or Monthly Costs by Linked Account.

A good way to look at the AWS Billing Console tools is as a way to map costs to the AWS parameters, like AWS Service, Linked Accounts, Regions, or RI usage. If you have a good idea of exactly how those parameters relate to your business, then you can use them to help gain valuable insights.

That being said, many businesses want to dig deeper into their AWS data and map it more directly to their company structure. That's where the AWS Cost and Usage Reports (CUR) come in. CUR files are the most comprehensive source for data about your AWS account. In your quest to fully manage your AWS spend, CUR is your ultimate tool and source of truth for cloud usage.

TIP

Do you use another provider, such as Azure or Google Cloud Platform, in addition to AWS? The CUR file can be used with [a multi-cloud cost management platform](#) so you can see all of your cloud spend in one place.

About the CUR File

The central promise of AWS is that you only pay for what you use, measured down to the second. The CUR file is where they prove that promise, providing detailed usage data that's incredibly granular. All told, a CUR file for a mature AWS infrastructure can contain hundreds of millions of lines.

As you can imagine, managing and storing a file this massive is a little trickier than just downloading a CSV file every now and then — especially since AWS updates the CUR many times a day. To help make storage and updates easier, part of setting up your account for CUR includes creating an S3 bucket to hold the CUR files. Accessing the CUR is then a matter of transferring the data from the S3 bucket, ideally to a tool that can make full use of the data. From the S3 bucket, the file can be downloaded, transferred into services like RedShift or directly queried using Amazon Athena.

To illustrate this setup, let's look at how CUR relates to Cloudability, where it's the source of truth for AWS accounts. When setting up accounts, we connect directly to the S3 bucket containing the CUR files. Whenever AWS updates the CUR file, we update the data in our platform.

The CUR file is the most complete data that you can possibly get about your AWS cost and usage, which also means it's the most versatile and flexible format for your data. The raw nature of CUR data allows you to normalize, analyze and map that data to your company's structure, which means you'll be able to get more dependable insights and make stronger business decisions about your AWS cloud.

Setting Up Your Account for CUR

Before you create the CUR, you need to make sure you have an S3 bucket specified. You can either use an existing one or create a new one, but the bucket should have a resource-based permission policy to make sure AWS can write files to it.

Once you have the S3 bucket created, go to the Reports section of your AWS Billing and Cost Management Console to create a report. You can get more details about the process in the [AWS documentation](#), but here are a few general recommendations for report parameters:

- Select *Include Resource IDs* under Additional Report Details for full granularity.
- Choose *Hourly* for Report Granularity to make sure you have as much data as possible.

DID YOU KNOW?

AWS used to use the Detailed Billing Report (DBR) instead of CUR. However, the DBR is now deprecated and [will be unavailable at a later date](#). If you're still using DBR, then it's a good idea to switch over to CUR.

PART 2:

Consolidated Billing

Consolidated Billing is part of AWS Organizations and is crucial for keeping large companies organized. Each Consolidated Billing organization has a single Master Account that is responsible for paying the bills of the Member Accounts.

There are many benefits to Consolidated Billing, especially if you have an enterprise-level cloud infrastructure:

- **One Bill:** Instead of having to gather multiple bills, all of your cost and usage data is gathered in a single source of truth.
- **Combined Usage:** Several services have volume discounts, such as Amazon S3. Consolidated Billing combines the usage of all accounts to take advantage of these discounts.

When you have Consolidated Billing set up, the Master Account is the only one with access to the CUR file. The CUR file will have the data from all of the Member Accounts, with each item tied to a specific account. You can also generate CUR files for specific Member Accounts.

TIP

Not sure where to begin or what the best setup is for your company? Don't worry, you're not alone. Our [Cloud Financial Management Services](#) can help you make sense of your billing and figure out the right AWS billing setup that fits your needs.

PART 3:

AWS Billing Data

What Can You Do with AWS Billing Data?

All right, you've got your billing accounts set up how you want them and your CUR file is stored in an S3 bucket. So now what do you do with those millions of lines of data?

The short answer is: completely understand your cloud. The CUR file gives you access to the raw cost and usage data for your cloud, which means you can move that data around to provide in-depth insights, recommendations and more.

No matter what you do with the data, the goal has to be to map your cost and usage to your company's structure, processes and infrastructure. Let's take a look at what the data enables you to do and how you can use it to make key business decisions.

Use Tags to Tie Cost and Usage to Your Business Structure

One of the most useful tools for tracking AWS usage is tagging. These key/value combinations can be applied to resources, then used to sort those resources to match your business. A common example is to apply tags by development environment, such as testing, production, qa, etc. They can also be used for teams, departments or even specific developers.


When you have all of your tagging set up correctly, you can get an incredible amount of visibility into your cloud cost usage. Best of all, the tagging maps it to your business structure, so you'll be able to instantly see how your cloud is being used.

TIP

The more tag coverage you have, the more you know about your billing and usage. But not everything can be tagged in AWS. Instead of looking at the tag coverage percentage, you should ask yourself, "Is everything tagged that could be tagged?"

Tag Explorer

A global view of coverage for all your tags and accounts groups. Learn more in our [Knowledge Base](#). Update your [tag mappings](#) or [account groups](#).



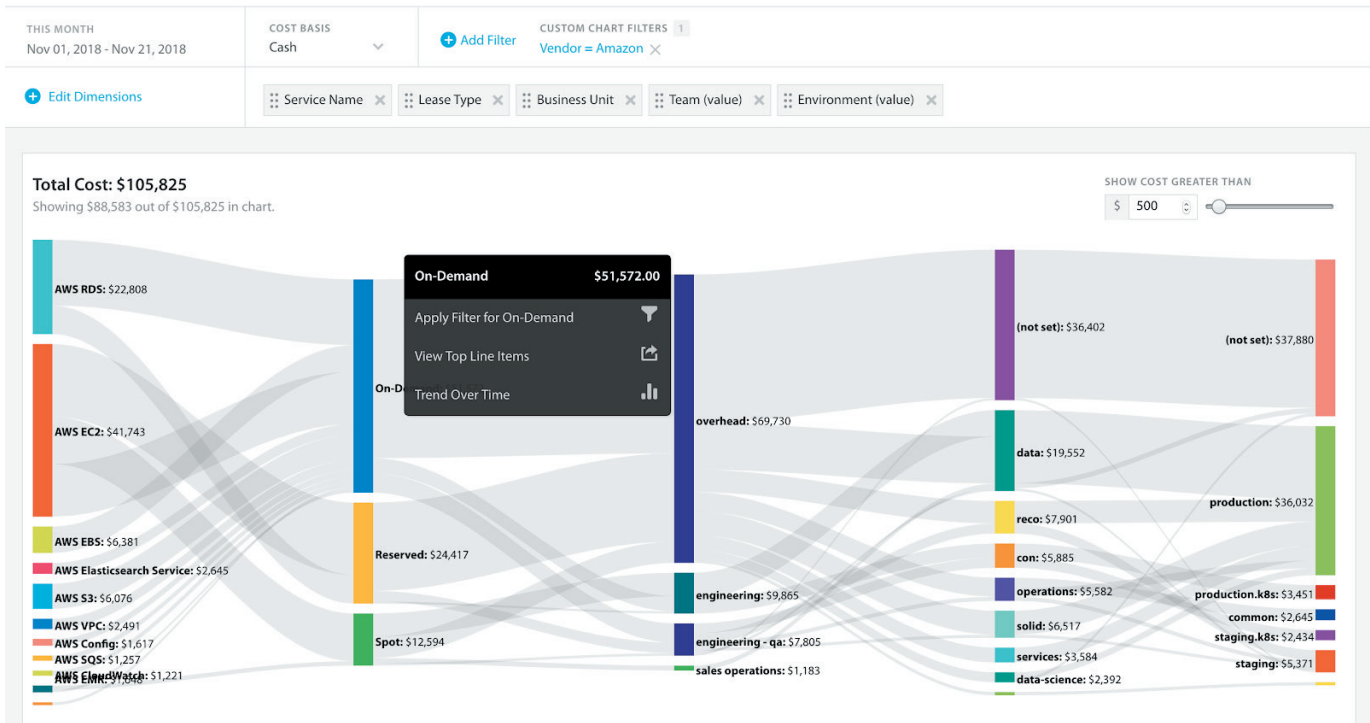
Cloudability's Tag Explorer shows you the distribution of tags visually and in table form. By using tags, you can instantly see costs and usage sorted by tags, including any untagged resources.

Tags can also be used to tie different data sets together, revealing relationships that aren't immediately apparent when looking at basic usage numbers. For example, distribution of costs by AWS service is a basic report. So is the split of spend between Reserved, Spot and On-Demand prices. But how do you map those costs to specific business units, teams and environments?

When used correctly, tags can allow you to look at the same data several different ways at the same time, giving you actionable insights into your cloud usage.

TrueCost Explorer

Explore your billing data visually and answer questions about how your usage translates into costs. [Learn More about TrueCost Explorer in our Knowledge Base.](#)



The Cloudability True Cost Explorer uses a next-level pivot table presentation to show the relationships between different dimensions of the same data.

The True Cost Explorer screenshot above shows you the kind of insights that tags can help you uncover. By looking at the data this way, you can see that a large percentage of RDS and EC2 is being purchased at the On-Demand rate — to the tune of \$41,743 total spend. You can also see who's using those On-Demand resources. Some are being used by Engineering and Engineering-qa, which makes sense in a production environment. But a large amount is being tagged as Overhead. In fact, most of the Overhead spend is at On-Demand rates instead of the cheaper Reserved or Spot pricing. Most importantly, you can also see that a large percentage of the On-Demand Overhead costs aren't tagged for a specific team or environment.

And that insight opens the door to lowering costs. By tracking down those untagged resources, you can figure out which team is using them, or if they're being used at all. You can also figure out if their use can be covered by Reservation pricing or if there's a way Spot Instances can be used. Since a solid RI and Spot instance strategy can save you 30-50% on your costs, you could generate \$12,500 to \$21,000 in savings over a three-week period — or \$215,000 to \$365,000 annually.

And that's only the tip of the iceberg. At Cloudability, we've found many ways to maximize tags and add new dimensions to the data with them. Our Tag Mapping feature accounts for typos (such as pooling Environment, environment, environiment, and prod-environ under a single environment tag), while Account Groups can create buckets for untaggable assets.

Tags are an invaluable tool for seeing what's really going on in your AWS Cloud. By building a strong tagging system and integrating it into AWS, you'll be able to supercharge the data you get in your CUR file and magnify the insights you can get from it.

TIP

To find out more about tags and building out a tagging strategy, check out our [AWS Tagging Strategy Best Practices](#) e-book.

Make the Most Out of Reservation Pricing

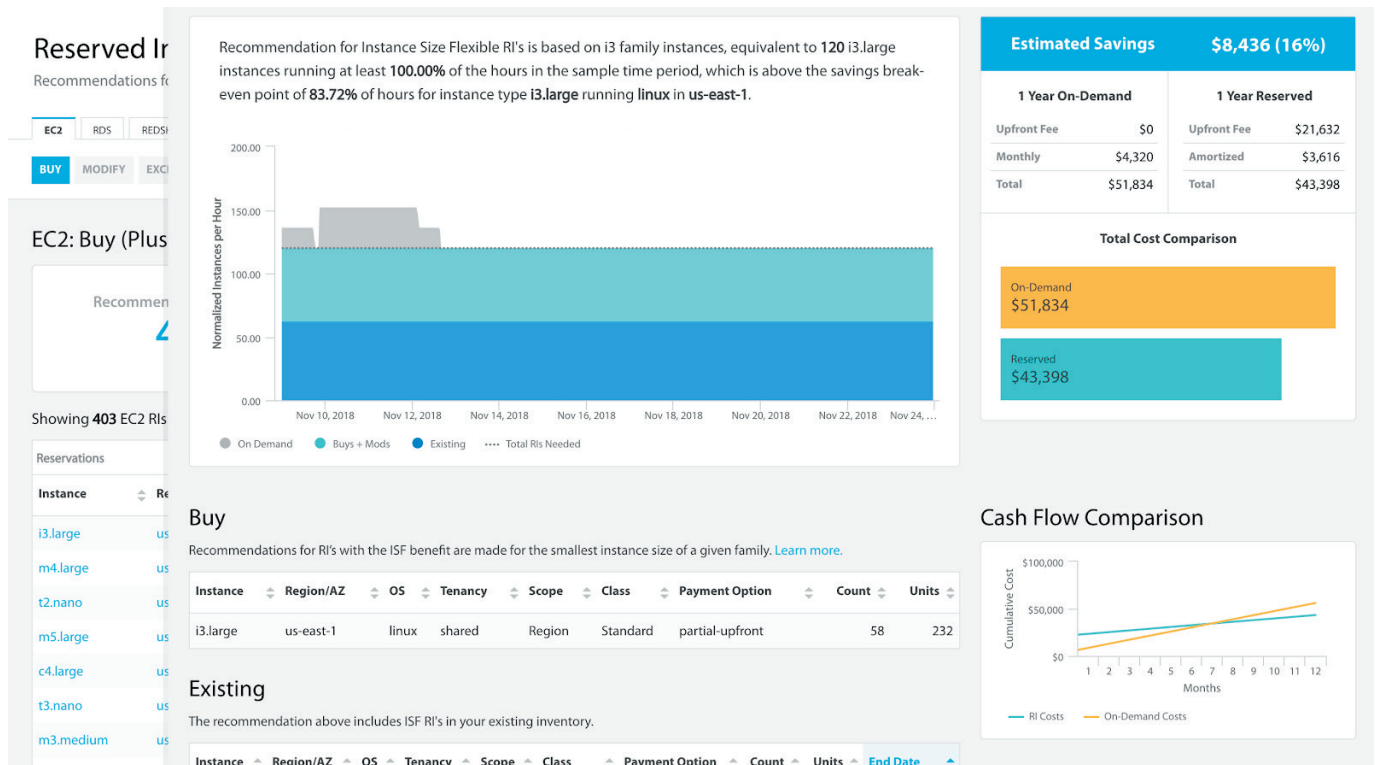
Some AWS services offer discounts for reserving a certain amount of usage over the course of a one-year or three-year period, including EC2, RDS, ElastiCache, Redshift and DynamoDB. If you know that you're going to have enough usage during that time period to cover the costs, then it's well worth the commitment. For example, Reserved Instances (the EC2 version of Reservations) are discounted up to 70% depending on the terms, parameters and payment options.

The trick is figuring out if you are going to use a Reservation enough to take advantage of the discount. An m5.large Standard RI with No Upfront in US West (Oregon) costs \$534.36 over the course of a full one-year term, an effective rate of \$0.061/hr. On-Demand for the same instance has an hourly rate of \$0.096, so you'd spend that same amount after using it for 5,566.25 hours, or about 64% of the time. If you want to make the commitment worthwhile, then you need to have the RI apply to a running m5.large instance at least 65% of the time, such as the instances behind your website or the ones behind key features that need 24/7 availability.

TIP

RIs can be one of the most powerful saving tools for AWS, but they can also be very complicated. Find out more about RIs with our [Complete Guide to Saving with AWS Reserved Instances](#).

Fortunately, all of the data you need to take advantage of RIs is located in your CUR file. RIs only apply to instances that fit specific parameters, including the family, size and region, but there are a lot of variables within those parameters. For example, an m5.large RI could be applied to one m5.large instance for 15 minutes, then another m5.large instance for the rest of the hour when the first is turned off. And if the RI has instance size flexibility, it could be applied to multiple smaller instances, like an m5.4xlarge RI applying to eight m5.large instances. Finding out whether you should purchase an RI involves comparing your usage data across multiple instances to make sure you'll get your money's worth.



Using CUR data, Cloudability's RI Planner makes recommendations for RI purchases that include how much you can save — and where the break-even point is for costs.

The larger your infrastructure gets, the harder it will be to uncover any RI savings potential. At the same time, larger infrastructures are more likely to maximize savings from RIs. The greater number of instances provides a greater chance for RIs to reach high utilization levels, and the bigger spend means the savings will add up faster.

Discovering potential RI savings gets even more complicated when you throw Convertible RIs into the mix. Like the name suggests, these RIs can be converted into other RIs with a similar price point. The flexibility means you're not locked into a specific RI type for the full term, and you can change your RI portfolio as your needs change. It also means that you have a lot more variables in play when finding savings opportunities.

By comparing the usage in your CUR file to the AWS pricing and RI data, you can create recommendations for Convertible RIs that will save you money without negatively affecting performance.

Reserved Instance Planner
Recommendations for buying new RIs, modifying current RIs for better coverage, and identifying underutilized RIs.

EC2 | RDS | REDSHIFT | ELASTICACHE | DYNAMODB **BETA**

BUY | MODIFY | **EXCHANGE** | UNDERUTILIZED

ACCOUNT: AWS Consolidated Master | RI OPTIONS: Standard / 1 Year

Convertible EC2 RIs

Reservation ID	Account Name	Instance	Project
	AWS Consolidated Master	r4.large	\$1,008.84
	AWS Consolidated Master	r3.large	\$1,092.00

Payment Details

CLASS	convertible	START DATE	Nov 07, 2018
PAYMENT OPTION	no-upfront	END DATE	Sep 07, 2019
TERM	one-year		

RI Details

INSTANCE TYPE	r4.large	COUNT	55
REGION/AZ	us-east-1	SCOPE	Availability Zone
TENANCY	shared	OS	linux

Savings Overview

By performing the steps below, you will have **maximized your utilization and savings** for this reservation

TOTAL SAVINGS
\$1,008.84
Over Remaining Reservation Term

STEP 1

Modify your reservation

Before you exchange your RIs, you need to modify your original reservation.

Instance	Region/AZ	OS	Tenancy	Term	Payment Option	Count
r4.large	us-east-1	linux	shared	one-year	no-upfront	48 ⓘ
r4.large	us-east-1	linux	shared	one-year	no-upfront	7

STEP 2

Exchange your RIs

After modifying your RI into pieces that can be exchanged, you'll want to exchange them.

Instance	Region/AZ	OS	Tenancy	Term	Payment Option	Count
<input checked="" type="radio"/> r4.large	us-east-1	linux	shared	one-year	no-upfront	7

EXCHANGE DETAILS

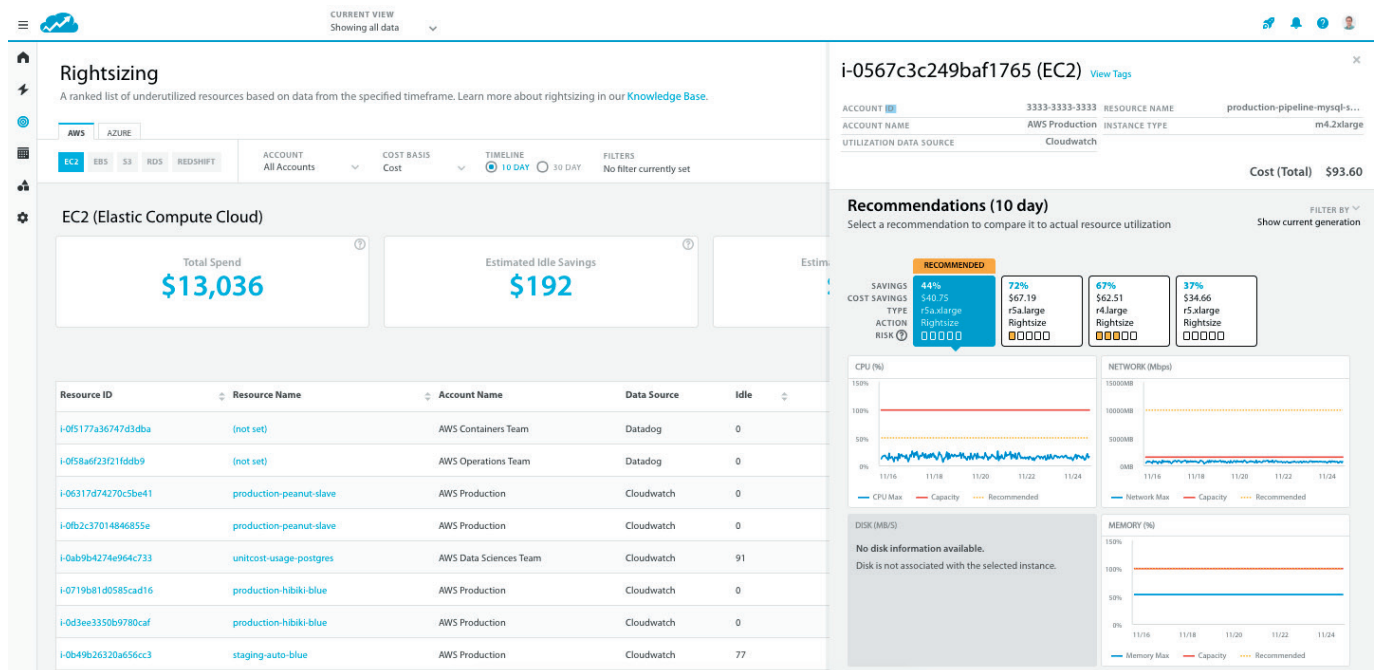
This example from the RI Planner shows how converting RIs can lead to substantial savings, demonstrating the power of using CUR data for creating RI recommendations.

Taking full advantage of RIs will mean different things to different companies depending on their available funds for investing in RIs, how changeable their cloud usage is and more. Whatever RI strategy you choose, your AWS billing data is essential to helping you reach your goals.

Optimize Current Cloud Usage

Data centers often required an overly cautious approach to provisioning. After all, it was better to invest a little more now to get extra capacity than to risk failure later on. Provisioning in the cloud requires a completely different approach. AWS makes it easy to spin up additional resources when needed, enabling you to optimize your current resource use to exactly what you need, then expand as your cloud scales up.

Using CUR data, you can track your actual usage, but that's only the beginning. Your billing and usage data can be integrated with utilization data using CloudWatch. Once you have that data, it can be cross-referenced against AWS instance data to create optimization recommendations.



There's always a risk of impacting performance when adjusting your provisioning. But the right analysis, like in our Rightsizing feature, can help you find the perfect option to minimize risk and maximize savings.

Optimization can have a significant impact on your cloud spend. When done correctly, you'll be able to free up valuable capital without negatively affecting the performance of your cloud. We've seen companies get the same performance from their cloud for 20% less just through optimization efforts. More often than not, that money is reinvested into new cloud initiatives that push the company forward.

All that being said, it is essential that you're able to have confidence in your optimization actions. Cut back too much and you can experience resource constraints that can, at best, slow down your applications and, at worst, crash your whole system. This is why it's so essential to have complete access to all of your utilization, billing and usage data for building data-driven predictions.

Fortunately, you have more than enough data to get the confidence you need. As an example, our Rightsizing pulls potentially hundreds of millions of lines in the CUR file to generate recommendations based on the past 10 or 30 days.

Optimization only becomes more vital the larger an infrastructure grows, and your billing data is essential to optimizing with confidence.

Next Steps

AWS billing data is a key part of getting the most from your cloud. No matter how your cloud is structured, you can use your cost and usage data to help you confidently make crucial decisions. It's really a matter of making sure you tie your cost and usage data to your organization's goals, processes and structures.

So your next step is to figure out the best way to tie the two together. Generally, we've seen companies start with spreadsheets, then progress to the built-in AWS tools or home-grown tools as their cloud grows. As the cloud scales up, we see companies either investing heavily in their custom tools or using a cloud cost management platform designed to take advantage of everything CUR has to offer.

In the end, you have to find the right setup for your business and your cloud. AWS billing data and the CUR file are powerful tools to fuel your cloud management. Make sure you have the machinery in place to take full advantage of the data.

Get Your Cloud Under Control

Whether you're a cloud-native company moving quickly or an enterprise looking to migrate to the cloud, there's a complex journey ahead. Get the resources to learn more about building and managing a cost-efficient cloud.

cloudability.com/resources

