

The Smart Way to Cloud: Your AWS Cost Optimization Hand book



CONTENTS

1 INTRODUCTION

2 ASSESSING YOUR CLOUD COSTS

3 COST OPTIMIZATION STRATEGIES AND BEST PRACTICES

4 NEXT STEPS IN YOUR COST OPTIMIZATION JOURNEY

The Smart Way to Cloud: Your AWS Cost
Optimization Handbook

INTRODUCTION

Chapter 1

www.cloudelligent.com

Common Challenges in Managing Cloud Costs

Welcome to the cost-conscious corner of the cloud!

A fall in expenditure owing to economic uncertainty has not lessened the importance of innovation and increased income creation for businesses. And cloud technology constitutes a catalyst for innovation. Driven by the rapid migration to the cloud and the complexity of cloud use and cost management, companies are struggling to control cloud spending and waste. Economic downturns, possible recessions, business cost-cutting, and revenue-boosting worsen this challenge. Cloud infrastructure management requires cost monitoring, resource optimization, and expenditure forecasting.



After surveying 753 global business and tech leaders, Flexera reveals a notable growth in cloud adoption, with 63% of respondents in 2022 indicating “heavy” use, up from 59% the previous year. Most enterprise organizations spend over a million dollars annually on cloud computing, with the prevalent expenditure for companies with over 1,000 employees ranging between \$2.4–\$6m. For SMEs, the common spend is \$600k–\$1.2m

“ However, there’s notable inefficiency in spending; on average, **32%** cloud expenditure is wasted, and projects typically run **13%** over budget. ”

By 2025, over 51% of enterprise IT expenditure is projected to be allocated to public cloud computing, in contrast to the 41% recorded in 2022, as forecasted by Gartner. In turn, this significant shift to the cloud represents a considerable proportion of IT budgets, prompting the pertinent question: Is your cloud spending optimized to its full potential? Given the increasing reliance on cloud platforms, it is crucial to discern the proportion of optimized cloud expenditure.

Recent Patterns and Trends in Cloud Costs

Last year, there was a noticeable uptrend in the pricing of cloud services, with virtually all providers augmenting their rates to differing extents. Brian Olsavsky, Amazon’s chief financial officer, observed that AWS “customers of all sizes in all industries” continuously search for avenues to save on costs. “Customers continue to evaluate ways to optimize their cloud spending in response to these tough economic conditions.” In a similar vein, the Uptime Institute highlights that the “cloud development projects are no different from many others and are likely to be postponed or deprioritized due to rising costs, skill shortages, and global uncertainty,” the report says.



94%
 acknowledge a rise in their cloud expenses (Virtana, 2023). More than half of these leaders, about 54%, have noted that the growth in storage costs is outpacing the increases in other cloud-related expenses

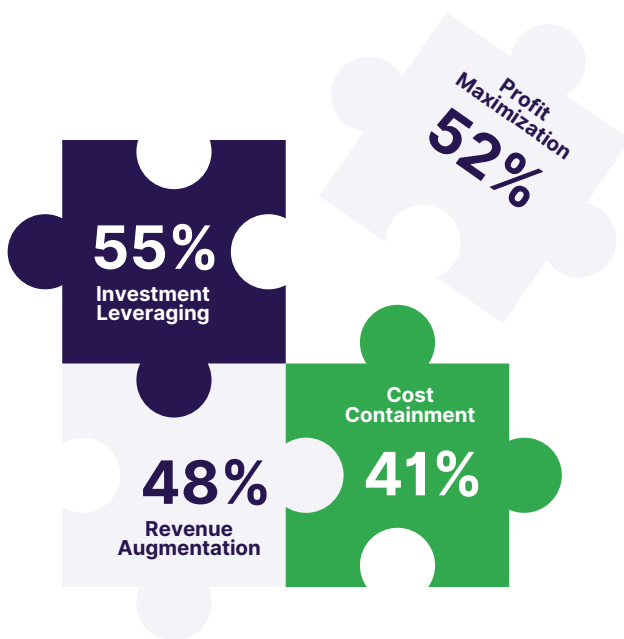
In this context, most IT leaders, approximately 94%, acknowledge a rise in their cloud expenses (Virtana, 2023). More than half of these leaders, about 54%, have noted that the growth in storage costs is outpacing the increases in other cloud-related expenses. Egress fees, often lurking in the shadows of cloud computing expense sheets, are more than mere line items; they accumulate rapidly, often catching businesses off guard - these unforeseen expenses not only lead firms to shell out considerably more than their projections but also constrict their agility to transition across cloud ecosystems.

In collaboration with 451 Research, AWS surveyed 1,000 IT decision-makers from various industries and 11 countries. Noteworthy findings include 95% of respondents believing that cloud services reduce Total Cost of Ownership (TCO) compared to on-premises solutions, with 72% of current public cloud users planning increased spending in the coming year to enhance productivity and revenue. Long-term cloud users with over 4 years of cloud usage are more likely to implement best practices and achieve significant unit cost savings of over 60%.

According to Liftr Insights data, AWS's average prices for on-demand compute instances have surged 23.0% over the past year. At the same time, AWS has raised its average prices consistently since 2019, and each year, the rate of increase has been steeper than the previous one.

Leveraging Cloud Cost Management to Boost Your Bottom Line

IBM distills the essence of cloud cost management, which “means knowing what your cloud operations cost and making intelligent adjustments so you can control cloud costs without compromising performance.” In this vein, S&P Global Market Intelligence explores four dimensions where cloud cost management impacts the business bottom line:



- **Investment Leveraging:** 55% of businesses have amplified their cloud investments, hinting at a larger strategic pivot.
- **Profit Maximization:** Over 52% have seen a tangible uptick in profitability, attributing this growth to the efficacy of cost management practices.
- **Revenue Augmentation:** Nearly half, at 48%, have experienced revenue growth, underscoring the potential for top-line impact.
- **Cost Containment:** For 41%, cloud financial strategies were critical for maintaining cost equilibrium, especially during volatile times.

Strategic, well-thought-out AWS cost management strategies do more than just cut costs; they redefine value, ensuring peak performance even as expenses dwindle. The resultant savings can be strategically reinvested, fostering business expansion and innovation. In addition, cost management clears performance bottlenecks.

Through this in-depth eBook, we go over all the details of managing AWS cloud resources. To make sure you get the most out of every dollar you spend, our goal is to give you a complete picture of how to use your cloud resources best. We talk about ways to cut costs while keeping your infrastructure's good performance and ability to grow. With the information in this guide, you can make better choices that will help your business succeed in the constantly changing digital world. As such, this guide is meant to help you learn how to optimize AWS costs and speed, no matter how much experience you have with the cloud or how new you are to it.

The Smart Way to Cloud: Your AWS Cost
Optimization Handbook

SURE ASSES

Assessing Your Cloud Costs

Chapter 2

ANALYS

www.cloudelligent.com

Demystifying your cloud bill is the first step on your cost-optimization journey.

In this section, we will guide you through the intricacies of understanding and analyzing your current AWS spend and lay the foundation for effective cost optimization strategies. This chapter breaks down the components of your cloud costs, spanning from compute and storage to the myriad services within the AWS ecosystem.

Additionally, we shed light on some dedicated tools for cost tracking and analysis. From native AWS tools such as AWS Cost Explorer to strategic third-party recommendations, we equip you with a robust toolkit to navigate your cloud expenses with precision and confidence.

Understanding the challenges of cloud cost optimization is key, especially with the intricate nature of cost billing. In fact, your AWS bill can contain hundreds and thousands of rows of extensive datasets, making it tough to efficiently manage costs. How do we tackle this?

AWS CUR 2.0: Understanding Your AWS Cost Structure

AWS Cost and Usage Reports (CUR 2.0) offers a powerful way to deeply understand and analyze your AWS spend. It provides the most comprehensive set of cost and usage data available for your account and empowers you to optimize cloud costs and avoid surprises on your bill.

Here are the basics of AWS cost structure and how CUR helps you make sense of it.

These provide the core information for each usage instance, including:

Line Items

- **Product:** Service used (e.g., Amazon EC2, Amazon S3)
- **Usage Type:** Specific action performed within the service (e.g., instance hours, data transfer)
- **Operation:** Specific API call or action that generated the usage
- **Cost:** Estimated charges associated with the usage
- **Resource ID:** Optional identifier for the specific resource used (e.g., EC2 instance ID, S3 bucket name)

Billing Details

Information about your billing and account, such as:

- **Billing Period:** Timeframe covered by the report (e.g., monthly)
- **Currency:** Used for cost calculations
- **Subscription ID:** Unique identifier for your AWS account

Identity Details

Information about the resources and users consuming the services, including:

- **Linked Account:** If billing for another account within your organization
- **Region:** AWS region where the usage occurred
- **Service:** Service that generated the usage
- **Resource Group:** AWS Resource Group associated with the usage
- **Tags:** User-defined tags attached to resources for categorization



A key visual for AWS CUR Report

	M	N	O	P	R	S	T
1	lineltem/ProductCode	lineltem/UsageType	lineltem/Operation	lineltem/AvailabilityZone	lineltem/UsageAmount	lineltem/CurrencyCode	lineltem/LineltemDescription
2	AmazonEC2	CW:AlarmMonitorUsage	Unknown		0.00134409	USD	\$0.00 per alarm-month - first 10 alarms
3	AmazonS3	Requests-Tier1	ListAllMyBuckets		2	USD	\$0.00 per request - PUT, COPY, POST, or LIST requests under the monthly global free tier
4	AmazonEC2	CW:AlarmMonitorUsage	Unknown		0.00134409	USD	\$0.00 per alarm-month - first 10 alarms
5	AmazonEC2	APS2-EBS:VolumeUsage-gp2	CreateVolume-Gp2		0.01344086	USD	\$0.00 per GB-month of General Purpose (SSD) provisioned storage under monthly free tier
6	AmazonEC2	APS2-EBS:VolumeUsage-gp2	CreateVolume-Gp2		0.01344086	USD	\$0.00 per GB-month of General Purpose (SSD) provisioned storage under monthly free tier
7	AmazonEC2	USW2-BoxUsage:t2.micro	RunInstances:0002	us-west-2a	1	USD	\$0.00 per Windows t2.micro instance-hour (or partial hour) under monthly free tier
8	AmazonEC2	USW2-USE1-AWS-Out-Bytes	PublicIP-Out		0.00000174	USD	\$0.000 per GB - data transfer out under the monthly global free tier
9	AmazonEC2	USW2-USE1-AWS-In-Bytes	PublicIP-In		0.00000138	USD	\$0.00 per GB - US West (Oregon) data transfer from US East (Northern Virginia)
10	AmazonEC2	USW2-USW1-AWS-In-Bytes	PublicIP-In		0.00000149	USD	\$0.00 per GB - US West (Oregon) data transfer from US West (Northern California)
11	AmazonS3	Requests-Tier1	ListAllMyBuckets		2	USD	\$0.00 per request - PUT, COPY, POST, or LIST requests under the monthly global free tier
12	AmazonEC2	USW2-DataTransfer-Out-Bytes	RunInstances		0.00038144	USD	\$0.000 per GB - data transfer out under the monthly global free tier
13	AmazonEC2	USW2-USW1-AWS-Out-Bytes	PublicIP-Out		0.00000174	USD	\$0.000 per GB - data transfer out under the monthly global free tier
14	AmazonEC2	USW2-DataTransfer-In-Bytes	RunInstances		0.00030951	USD	\$0.000 per GB - data transfer in per month
15	AmazonEC2	USW2-BoxUsage:t2.micro	RunInstances:0002	us-west-2a	1	USD	\$0.00 per Windows t2.micro instance-hour (or partial hour) under monthly free tier
16	AmazonEC2	USW2-USW1-AWS-Out-Bytes	PublicIP-Out		0.00000349	USD	\$0.000 per GB - data transfer out under the monthly global free tier
17	AmazonEC2	USW2-USW1-AWS-In-Bytes	PublicIP-In		0.00000276	USD	\$0.00 per GB - US West (Oregon) data transfer from US West (Northern California)
18	AmazonEC2	APS2-EBS:VolumeUsage-gp2	CreateVolume-Gp2		0.01344086	USD	\$0.00 per GB-month of General Purpose (SSD) provisioned storage under monthly free tier
19	AmazonEC2	CW:AlarmMonitorUsage	Unknown		0.00134409	USD	\$0.00 per alarm-month - first 10 alarms
20	AmazonEC2	USW2-BoxUsage:t2.micro	RunInstances:0002	us-west-2a	1	USD	\$0.00 per Windows t2.micro instance-hour (or partial hour) under monthly free tier
21	AmazonEC2	USW2-DataTransfer-Regional-Bytes	PublicIP-Out		0.00000349	USD	\$0.000 per GB - regional data transfer under the monthly global free tier
22	AmazonEC2	USW2-DataTransfer-In-Bytes	RunInstances		0.00032071	USD	\$0.000 per GB - data transfer in per month
23	AmazonEC2	USW2-DataTransfer-Regional-Bytes	PublicIP-In		0.00000302	USD	\$0.000 per GB - regional data transfer under the monthly global free tier
24	AmazonEC2	USW2-USE1-AWS-Out-Bytes	PublicIP-Out		0.00000174	USD	\$0.000 per GB - data transfer out under the monthly global free tier
25	AmazonEC2	USW2-DataTransfer-Out-Bytes	RunInstances		0.00045736	USD	\$0.000 per GB - data transfer out under the monthly global free tier
26	AmazonEC2	USW2-DataTransfer-In-Bytes	RunInstances		0.00036737	USD	\$0.000 per GB - data transfer in per month
27	AmazonEC2	USW2-APN2-AWS-In-Bytes	PublicIP-In		0.00000005	USD	\$0.00 per GB - US West (Oregon) data transfer from Asia Pacific (Seoul)
28	AmazonEC2	USW2-APN2-AWS-Out-Bytes	PublicIP-Out		0.00000018	USD	\$0.000 per GB - data transfer out under the monthly global free tier
29	AmazonEC2	USW2-USE1-AWS-In-Bytes	PublicIP-In		0.00000153	USD	\$0.00 per GB - US West (Oregon) data transfer from US East (Northern Virginia)
30	AmazonEC2	USW2-DataTransfer-Out-Bytes	RunInstances		0.00039945	USD	\$0.000 per GB - data transfer out under the monthly global free tier
31	AmazonEC2	CW:AlarmMonitorUsage	Unknown		0.00134409	USD	\$0.00 per alarm-month - first 10 alarms

AWS CURs are essentially spreadsheets with extensive rows and columns. Despite containing all necessary data, extracting insights can become challenging due to the lack of visualization options. Therefore, we recommend using Amazon QuickSight, a serverless business intelligence service that can connect to your AWS CUR data and provide powerful visualizations, insights, and recommendations.

Identifying Key Cost Drivers on AWS

In line with the latest iterations from [AWS from February 2023](#), the costs associated with AWS services are primarily driven by three fundamental factors:

- Compute
- Storage
- Outbound data transfer

These cost drivers play a central role in determining your overall cloud expenditure. It's essential to understand that the specifics of these cost drivers may vary depending on the AWS resource and pricing model you opt for.

Compute Costs

- **Instance Types:** Choosing Amazon EC2 (Elastic Compute Cloud) instance types significantly impact compute costs. Different instance types offer varying CPU, memory, and specialized hardware levels, with corresponding price differences.
- **Usage Duration:** The longer you run instances, the higher the compute costs. AWS charges hourly or per second, depending on the instance type, so optimizing usage duration can save costs.

Storage Costs

- **Storage Types:** AWS offers a range of storage types, each with its associated cost structure. These include Amazon S3 (Simple Storage Service), Amazon EBS (Elastic Block Store), and more. The choice of storage type should align with your performance and cost requirements.
- **Data Volume:** The volume of data you store directly affects storage costs. Watch out for data retention policies and regularly assess and archive no longer-needed data.
- **Snapshots and Backups:** Creating frequent snapshots and backups can contribute to increased storage costs. Implement lifecycle policies to manage snapshots and backups efficiently.

Outbound Data Transfer Costs

- **Data Transfer Volume:** The volume of data transferred from AWS to the internet or between AWS regions influences outbound data transfer costs. Be aware of data transfer rates when designing your architecture.
- **Content Delivery:** Utilizing content delivery networks (CDNs) like Amazon CloudFront can reduce data transfer costs by caching and serving content closer to end-users.

Service-Specific Costs

It's important to note that some AWS services may have their specific cost drivers. For instance, data processing services like AWS Lambda may incur charges based on the number of requests or compute time.

Database Costs

Database costs can be a significant expense for any organization, particularly those with large or complex data sets. They can be driven by several factors such as the size and complexity of the database, the complexity of data queries, the frequency of backups and data transfers.

Network Costs

Network costs are another critical aspect of cloud spend that organizations should pay close attention to. These costs are influenced by data transfer volumes, bandwidth usage, and the geographical distribution of cloud resources.

Every AWS service has a price tag, from the computing power of EC2 instances to the storage breadth of S3 buckets. Yet, prices aren't one-size-fits-all. The region you're in, the data you transfer, and even the specific service type can swing costs. Without proper management and monitoring, expenses can pile up. As dynamic as these services are, so is the importance of managing them effectively.

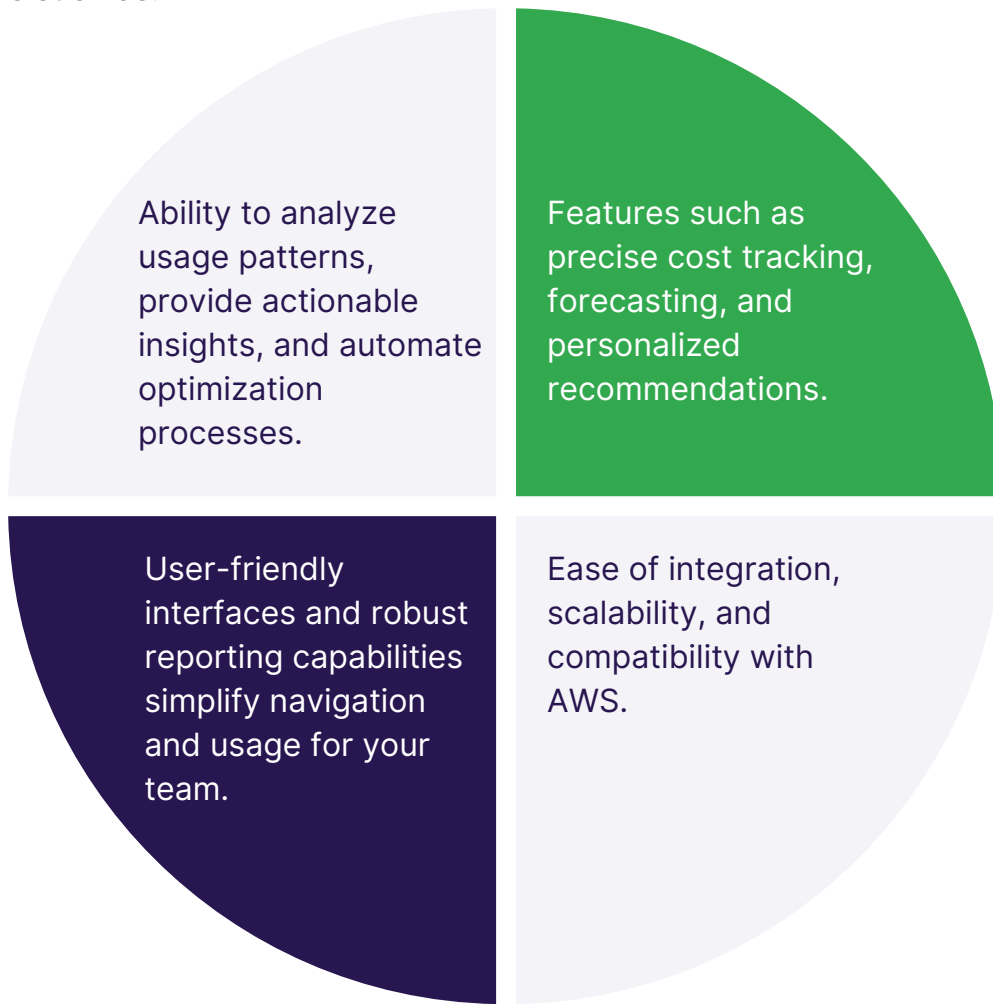
Now that we've pinpointed the key cost drivers within your AWS environment, it's time to equip ourselves with the right tools to tackle them head-on.

Selecting the Right Tools for Cost Optimization

So what exactly are Cloud Cost Optimization Tools?

A Cloud Cost Optimization Tool is an application or service designed to help businesses manage and reduce their cloud expenses effectively. These tools analyze usage patterns, identify inefficiencies, and provide recommendations to optimize resource allocation in cloud environments. They often offer insights into spending, automate cost control measures, and help organizations make informed decisions to maximize efficiency.

And how do you know which Cloud Cost Optimization Tool aligns with your business needs? Consider factors such as:



By choosing a tool that aligns with your requirements, you can pave the way for improved cost visibility, control, and overall efficiency.

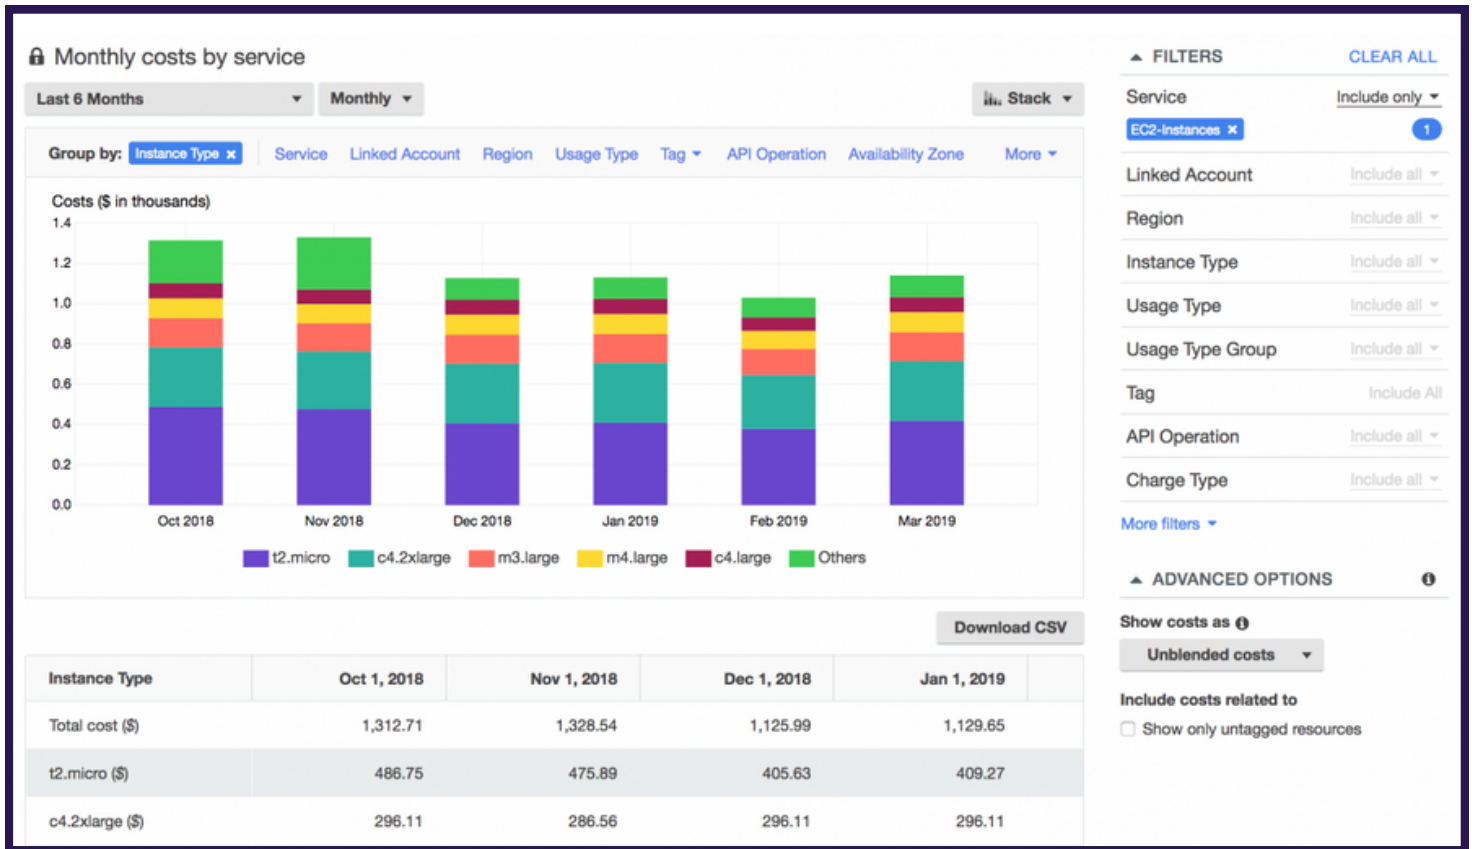
Native AWS Tools for Assessing and Managing Cloud Costs

AWS provides a set of powerful tools to help you assess, track, and manage your cloud expenditures effectively. These tools are designed to offer a holistic view of your AWS usage and empower you to make informed decisions and enhance cost control. Here's a closer look at some native AWS cost optimization services:

AWS Cost Explorer



A key Visual for AWS Cost Explorer displaying monthly costs by service.



This is your go-to tool for visualizing, understanding, and managing AWS costs and usage. It provides an intuitive interface for exploring historical data, forecasting future costs, and identifying cost-saving opportunities. You can leverage features such as:

- **Comprehensive Data Visualization:** Gain insights through intuitive and interactive visualizations that simplify the interpretation of complex cost data.
- **Historical Cost Analysis:** Explore your cost trends over time and enable informed decision-making based on historical usage patterns.
- **Future Forecasting:** Predict future spending based on historical usage which empowers you to budget wisely and avoid nasty surprises.
- **Customized Reports:** Generate tailored reports to focus on specific cost aspects and facilitate targeted analysis for different stakeholders.

AWS Cost Anomaly Detection



A key visual for anomalies detected by service and cost.

Anomalies detected (47) Info						
<input type="text" value="Find detected anomalies by property or value"/>				Last 90 days (all) ▼		< 1 2 3 4 5 >
Detection date ▼	Severity ▼	Duration	Service	Account ID	Total cost impact ▲	Assessment
2020-11-11	Low	14 days	Amazon Elastic Compute Cloud - Compute		\$3,348,575.96	Not submitted
2020-10-21	Low	6 days	Amazon Elastic Compute Cloud - Compute		\$1,046,571.62	Not submitted
2020-10-08	Low	8 days	Amazon Elastic Compute Cloud - Compute		\$953,814.27	Not submitted
2020-11-15	Low	3 days	Amazon Elastic Compute Cloud - Compute		\$898,037.87	Not submitted
2020-10-26	Low	3 days	Amazon Elastic Compute Cloud - Compute		\$669,155.71	Not submitted
2020-11-29	Low	1 day	Amazon Elastic Compute Cloud - Compute		\$390,791.80	Not submitted
2020-11-22	Low	1 day	Amazon Elastic Compute Cloud - Compute		\$346,158.89	Not submitted
2020-09-29	Low	1 day	Amazon Elastic Compute Cloud - Compute		\$70,607.67	Not submitted
2020-10-27	Low	4 days	Amazon Elastic Block Store		\$39,900.74	Not submitted
2020-10-23	High	3 days	Amazon Virtual Private Cloud		\$29,734.04	Not submitted

This is a neat feature within AWS Cost Explorer that leverages machine learning to identify and alert users about unusual spending patterns. It helps in detecting anomalies in cost and usage data and empowers your business to quickly respond to unexpected changes and potential cost inefficiencies. Some features include:

- **Proactive Anomaly Detection:** Uses machine learning to identify unusual spikes in your AWS spending compared to historical patterns.
- **Real-time Alerts:** Receive real-time notifications when anomalies are detected so your team can take immediate action.
- **Cost Trend Analysis:** Analyzes historical cost data to identify trends and anomalies over time.
- **Actionable Insights:** Receive recommendations on potential cost-saving actions for any identified anomalies.

AWS Trusted Advisor



A key visual for potential monthly savings.

Cost optimization

Choose a check name to see recommendations for ways to help save money for your AWS account. Trusted Advisor might recommend that you delete unused and idle resources, or use reserved capacity.

Overview

- Potential monthly savings: **\$7,082.26**
- Action recommended: **1**
- Investigation recommended: **14**
- No problems detected: **10**
- Checks with excluded items: **11**

Cost optimization checks

Filter by tag key [Learn more about using tags](#)

Tag Key Tag Value

Search by keyword [Info](#) Source View

Amazon Comprehend Underutilized Endpoints

Checks the throughput configuration of your endpoints.

Last updated: 2 hours ago

A tool that provides best practices and recommendations to optimize your AWS environment, including cost optimization suggestions. It helps identify potential cost savings opportunities and improve overall performance.

- **Automated Recommendations:** Identifies optimization opportunities across your AWS services, including security, performance, and cost.
- **Cost Optimization Insights:** Offers best practices and recommendations for optimizing costs across various AWS services.

AWS Budgets



A key visual displaying the current budgets.

AWS Budgets

Filter by budget name

All budgets (7) Cost budgets (5) Usage budgets (2) Reservation budgets (0)

Budget name	Budget type	Current	Budgeted	Forecasted	Current vs. budgeted	Forecasted vs. budgeted
Project Nemo Cost Budget	Cost	\$43.90	\$45.00	\$56.33	97.55%	125.17%
Eastern US Regional Budget	Cost	\$85.21	\$100.00	\$125.28	85.21%	125.28%
Total Monthly Cost Budget	Cost	\$141.50	\$175.00	\$187.00	80.86%	106.86%
Total EC2 Cost Budget	Cost	\$136.90	\$200.00	\$195.21	68.45%	97.61%

This enables you to set custom cost and usage budgets that alert them when thresholds are exceeded. AWS Budgets is instrumental in proactively managing costs by providing real-time notifications and helping businesses stay within their defined budgetary limits.

- **Customizable Budgets:** You can set custom budgets based on various criteria such as services, linked accounts, and tags.
- **Budget Action Automation:** Trigger automatic actions (e.g., resource stop/termination) upon reaching specific budget thresholds.
- **Automated Monitoring and Alerts:** Receive notifications when your cloud spend approaches or exceeds predefined budget thresholds. You can then take immediate action, accordingly.
- **Customizable Reports:** Generate detailed reports on cost allocations, resource usage, and budget utilization.

Third-Party Solutions for AWS Cost Optimization

While native AWS tools equip you with a strong foundation for assessing and managing your cloud costs, consider them the launchpad - not the final destination. Third-party tools often bring a more diverse range of features, customization options, and integrations that may not be present in the native AWS offerings. These tools are developed by specialized vendors with a focus on advanced analytics, predictive modeling, and automation capabilities.

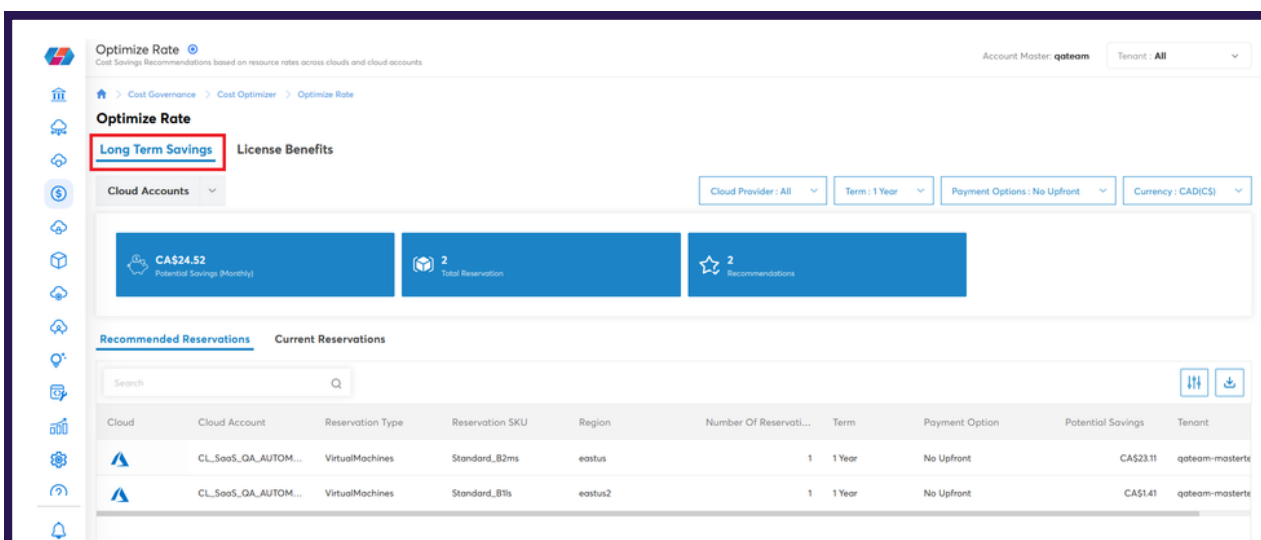
In this section, we'll recommend some amazing tools - a curated selection of third-party solutions designed to complement and amplify your cost optimization efforts. These platforms delve deeper into specific areas such as granular cost analysis, cloud spend forecasting, and even reservation management.

For Managing Cloud Costs and Optimizing Resources



Recommendation 1

CoreStack: AI-Powered Cloud Governance



The screenshot displays the 'Optimize Rate' interface in CoreStack. It features a sidebar with navigation icons and a main content area. A red box highlights the 'Long Term Savings' section. Below this, there are three summary cards: 'CA\$24.52 Potential Savings (Monthly)', '2 Total Reservations', and '2 Recommendations'. A table titled 'Recommended Reservations' is visible, with columns for Cloud, Cloud Account, Reservation Type, Reservation SKU, Region, Number Of Reservations, Term, Payment Option, Potential Savings, and Tenant.

Cloud	Cloud Account	Reservation Type	Reservation SKU	Region	Number Of Reservations	Term	Payment Option	Potential Savings	Tenant
🇺🇸	CL_SaaS_GA_AUTOM...	VirtualMachines	Standard_B2ms	eastus	1	1 Year	No Upfront	CA\$23.11	qateam-masterte
🇺🇸	CL_SaaS_GA_AUTOM...	VirtualMachines	Standard_B1ts	eastus2	1	1 Year	No Upfront	CA\$1.41	qateam-masterte

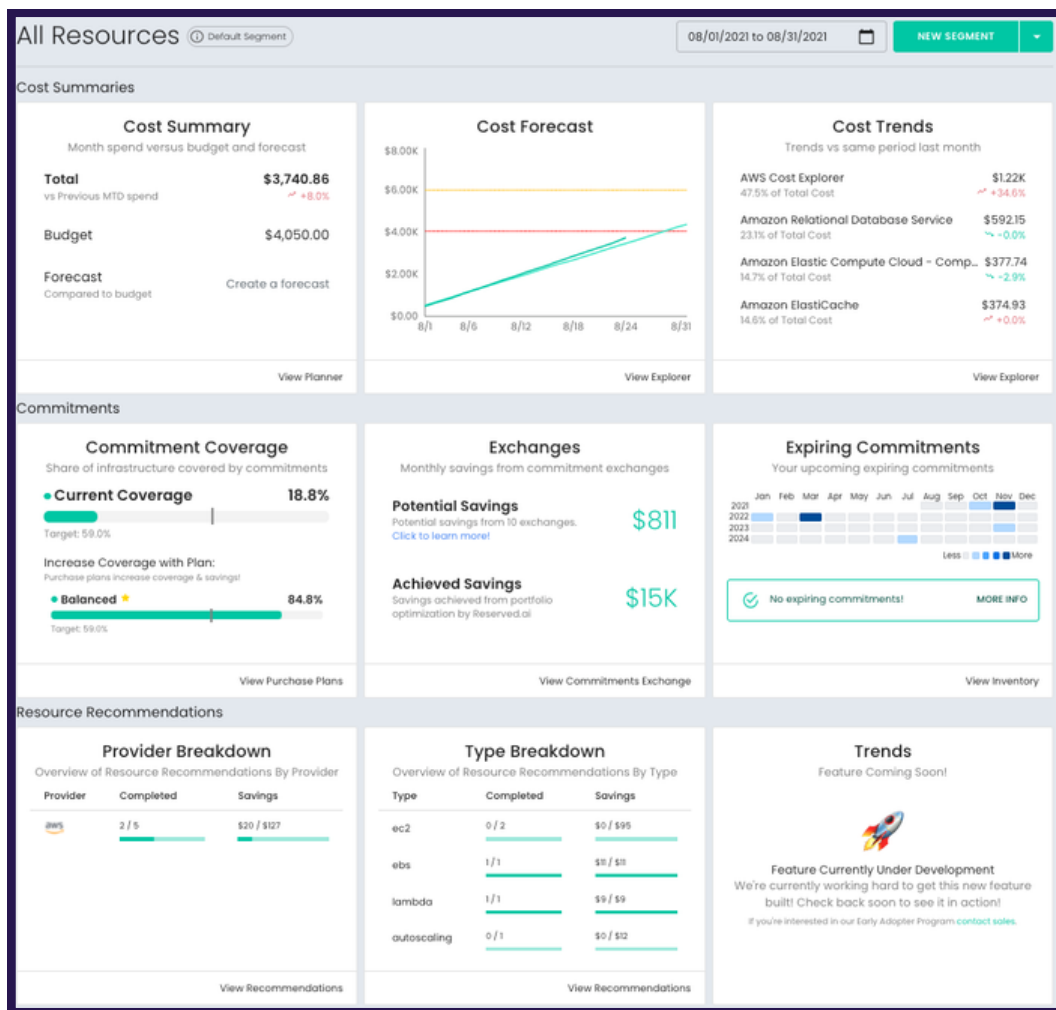
At Cloudeelligent, we highly recommend CoreStack. This AI-powered platform is a game-changer for businesses that offers better cost control, accurate predictions, and optimized cloud resource utilization. Plus, it promotes fiscal responsibility with real-time alerts and proactive monitoring. With CoreStack, we're confident in helping clients optimize their FinOps framework for maximum efficiency and savings.

For Reservation Management and Orchestration



Recommendation 2

Archera: Insure Commitments for RIs and Savings Plans



We are enthusiastic about Archera's capabilities in optimizing cloud resource allocation and expenditure management. It's a sophisticated solution that taps into advanced analytics and machine learning to dissect cloud usage patterns. With its insightful recommendations, Archera becomes your ally in forecasting, optimizing, and securing AWS commitments.

Archera's 30-Day alternatives to Reserved Instances and Savings Plans offer flexibility, ensuring coverage and savings without long-term contracts. Plus, access to 30-Day Guaranteed Reserved Instances (GRIs) and comparison tools makes cost-saving strategies a breeze.

The Smart Way to Cloud: Your AWS Cost
Optimization Handbook

Cost Optimization Strategies and Best Practices

Chapter 3


www.cloudelligent.com

AWS defines Cost Optimization as “the ability to run systems to deliver business value at the lowest price point.”

But achieving this ideal isn't as straightforward as it sounds. The real question becomes: does increased spending translate to increased value? Chapter 3 is your comprehensive guide to building a FinOps framework that ensures every dollar you put into the cloud brings you the maximum value.

Adopting the AWS Well-Architected Framework

The AWS Well-Architected Framework is a set of best practices to build secure, high-performing, resilient, and efficient cloud architectures. It helps businesses optimize their workloads by providing guidelines across six pillars, including the Cost Optimization Pillar.



6 Pillars of AWS Well-Architected Framework

By following the framework's recommendations, businesses can enhance their cloud architecture's efficiency and ensure that costs align with actual usage. Regularly assessing workloads based on the Cost Optimization Pillar enables ongoing improvements. This ultimately leads to reduced operational costs and improved financial management of AWS resources.

Quick Cost Optimization Wins You Can Adopt Immediately

Even though some strategies for cost optimization focus on the long-term and require diligent monitoring, there are a few actions that you can take immediately:

Analyze usage with S3 Analytics and shift data to cost-effective storage options like S3 IA using Life Cycle Policies or S3 Intelligent-Tiering.



Manage Amazon S3 Storage Costs

Utilize AWS Compute Optimizer for EC2 rightsizing recommendations. Employ AWS Instance Scheduler for auto-stopping and AWS Operations Conductor for auto-resizing instances.



Manage Amazon EC2 Costs

Critically review and modify EC2 Autoscaling Group configurations and scaling policies for cost-effectiveness without compromising performance.



Adjust EC2 Autoscaling Groups

Cost Optimization Design Principles



Building a FinOps Framework



Gaining Cost Transparency and Control



Choosing Cost-Optimized Resources



Continuously Optimizing Over Time



Best Practices for AWS Cost Optimization

Cloud cost optimization is dynamic and adapts to shifting application requirements and service and pricing options. Yet, it is feasible to optimize the cost of your cloud environment by adhering to the following best practices.

Building a FinOps Framework

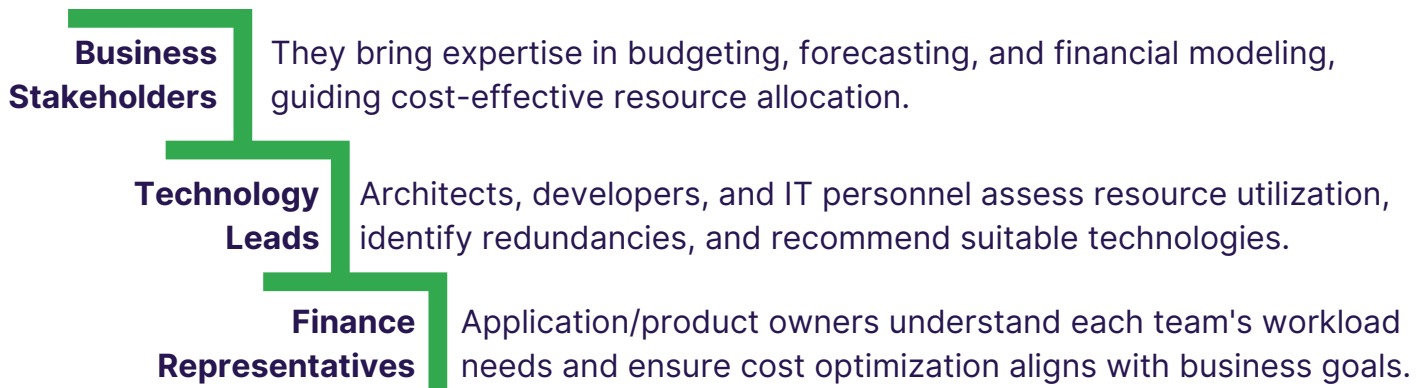
Managing cloud finance requires evolving your existing business processes to operate with cost transparency, control, planning, and optimization of your AWS environment. Adopting the FinOps Framework empowers multiple teams to effectively manage costs as your AWS usage scales.

We recommended the following FinOps strategies to optimize your financial management within AWS.

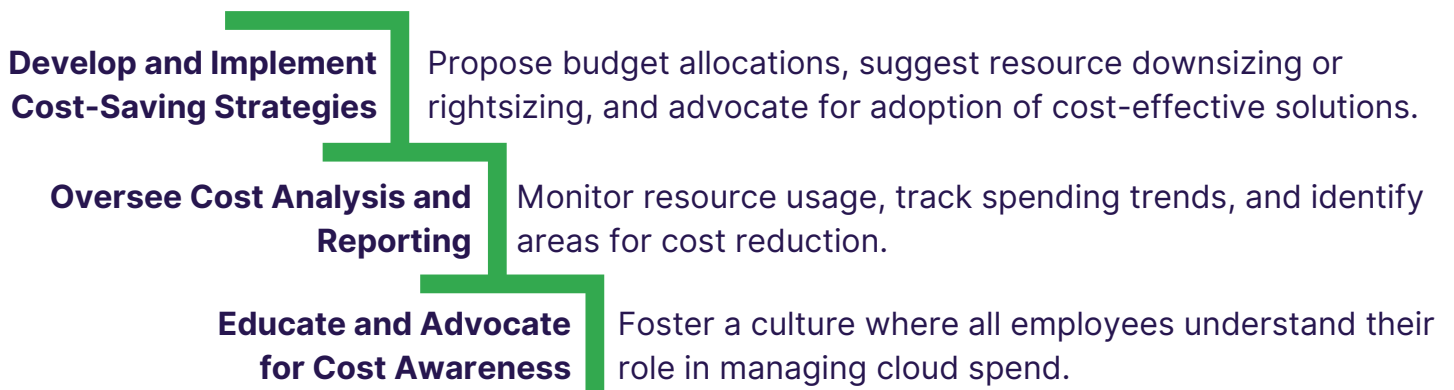
Take Ownership of Cost Optimization

Optimizing AWS costs isn't a solo act – it's a team effort driven by clear accountability and shared responsibility. But where do you begin? By establishing ownership, you create a framework for success and ensure everyone understands their role in achieving financial agility. Here are some steps you can implement:

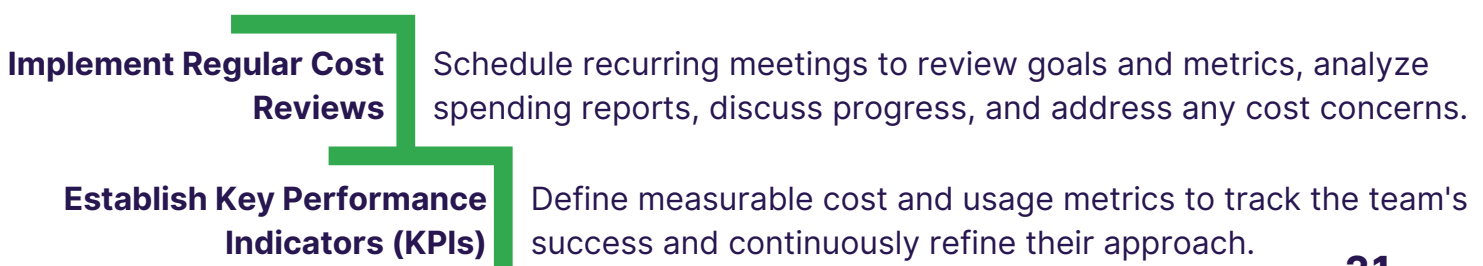
Establish a Dedicated Team: Assemble a team specifically tasked with overseeing and implementing cost optimization strategies. Therefore, you should ideally include:



Assign Responsibilities and Tasks: Clearly outline the specific tasks and duties of each team member to ensure a cohesive effort toward achieving your cost optimization objectives. Focus on these tasks:



Ensure Effective Ownership: By empowering your team with the necessary authority and resources, you can ensure their ownership translates into proactive action and tangible cost savings for your business.



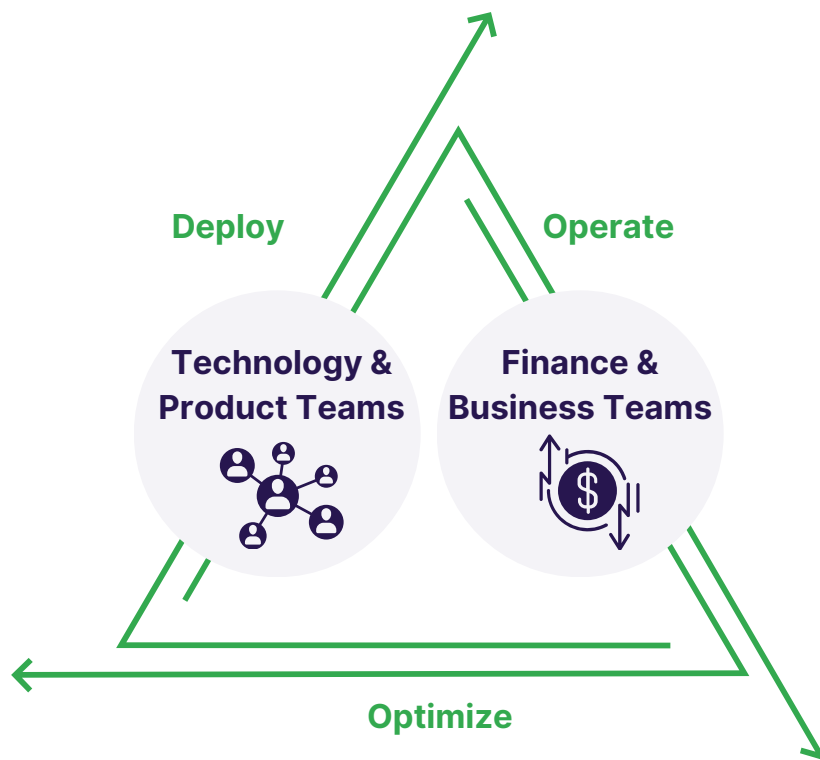


Facilitate Collaboration Between Finance and Technology

Keep the finance and tech teams in the loop during every step of your cloud journey. Facilitate regular meetings where teams delve into crucial topics such as business objectives and the current state of cost and usage.

To succeed, finance and technology must become equal partners in managing cloud spend.

- **Finance Leads:** CFOs, controllers, and planners understand cloud models and billing, pay-as-you-go pricing, and detailed usage data.
- **Technology Leads:** Product and application owners grasp financial constraints and business requirements to optimize deployment for desired goals.



Establish Cloud Budgets and Forecasts

AWS costs can be as dynamic as your workloads. They can increase with unexpected traffic or shrink during off-peak hours. This is where budgets and forecasts come into play. Let's break down the process of establishing effective controls:

Budgeting

- **Apply Granular Control:** Allocate budgets by department, project, or service which enable ownership and accountability.

- **Set Cost Guardrails:** Leverage AWS Budgets Alerts and AWS Cost Anomaly Detection to get alerts and notifications for preventing runaway costs.
- **Perform Regular Reviews:** Adjust budgets based on business goals and changing workloads with quarterly or monthly reviews.

Analyzing Historical Costs

- **Trend-Based Forecasting:** Utilize tools such as AWS Cost Explorer to analyze past usage patterns across services, accounts, and tags. Identify seasonal variations, growth trends, and anomalies.
- **Driver-Based forecasting:** Collaborate with business units to understand upcoming initiatives, marketing campaigns, or product launches. Quantify their potential impact on cloud resource consumption.



Cultivate a Culture of Cost Consciousness

Optimizing cloud costs isn't a one-person show; it's a cultural shift. Embedding cost awareness into your organization's DNA requires weaving it into existing processes. Here's what you can do:

- **Identify Cost-Impacting Processes:** Each team scans their workflows for anything that spins up AWS resources or shuts them down.
- **Cultivate a Cost-Aware Mindset:** Time to make cost awareness a team sport! Make sure all the stakeholders understand how their actions impact the cloud budget.
- **Integrate Cost Awareness into Processes:** Modify your processes to be cost-conscious and integrate cost awareness into training and incident management.



Gaining Cost Transparency and Control

Understanding your organization's costs and drivers is crucial for efficient cost management. We highly suggest you consider embracing a multi-faceted approach to gain insight into your usage and expenses. Follow these key best practices:

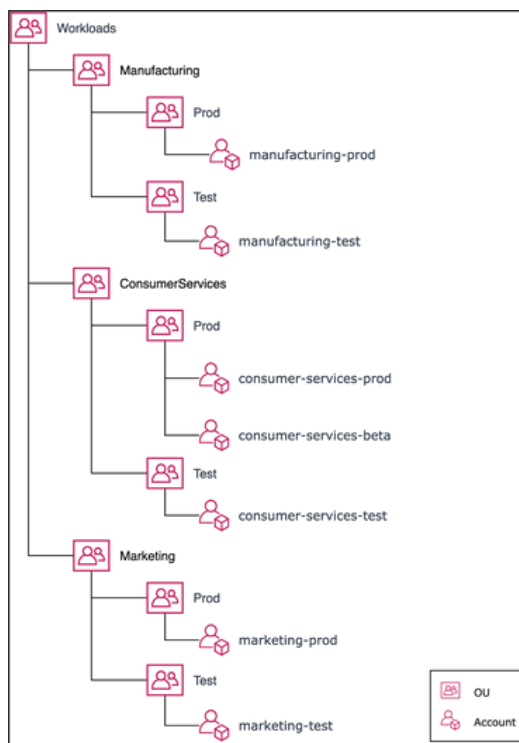


Implement a Robust Governance Framework

Think of governance as the traffic rules for how your business uses AWS resources. It helps to align your cloud strategy with your business objectives, ensure compliance with regulations and best practices, and optimize costs.

For effective cloud governance, here’s what you can implement:

- **Policies Based on Business Requirements:** Policies are the rules that specify what actions are allowed or denied on your AWS resources. Define policies for provisioning and decommissioning, usage limits, tagging, and actions to prevent runaway bills.
- **Goals and Targets:** Set both cost and usage goals for your workloads. Goals guide business direction, while targets ensure measurable outcomes.
- **Account Structures:** Design an account structure that reflects your organizational hierarchy and organize by project, departments, or workloads.



A key visual for an account structure by department.

- **Groups and Roles:** Assign granular access permissions to prevent unauthorized spending.
- **Cost Controls:** Leverage AWS Budgets for real-time alerts on spending deviations from policies. Utilize AWS Cost Anomaly Detection to mitigate any surprise costs and analyze root causes.
- **Project Lifecycles:** Implement a tracking mechanism for the entire project lifecycle. Automate resource termination at project completion to avoid lingering costs.



Establish Monitoring and Reporting Mechanisms

Knowledge is power! When it comes to cloud costs, visibility is your ultimate weapon. Monitoring cost and usage helps you collect and analyze data related to your cloud consumption and spend.

To achieve this, you can proceed as follows:

- **Configure Detailed Information Sources:** Leverage data sources such as AWS Cost and Usage Reports, AWS Cost Explorer, and AWS CloudWatch) to access the details of your cloud usage and cost.
- **Categorize Costs for Attribution:** These are the dimensions that help you group and allocate your cloud costs. You can identify categories such as project, department, or application to pinpoint where costs originate.
- **Define Organization Metrics:** Establish key performance indicators (KPIs) such as cost per user, cost per API call, or resource utilization rate for each workload. Track these metrics over time to identify trends and areas for improvement.
- **Leverage Billing and Cost Management Tools:** Refer to Chapter 2 of this handbook for more details on AWS native and third-party tools.



Decommission Unused AWS Resources

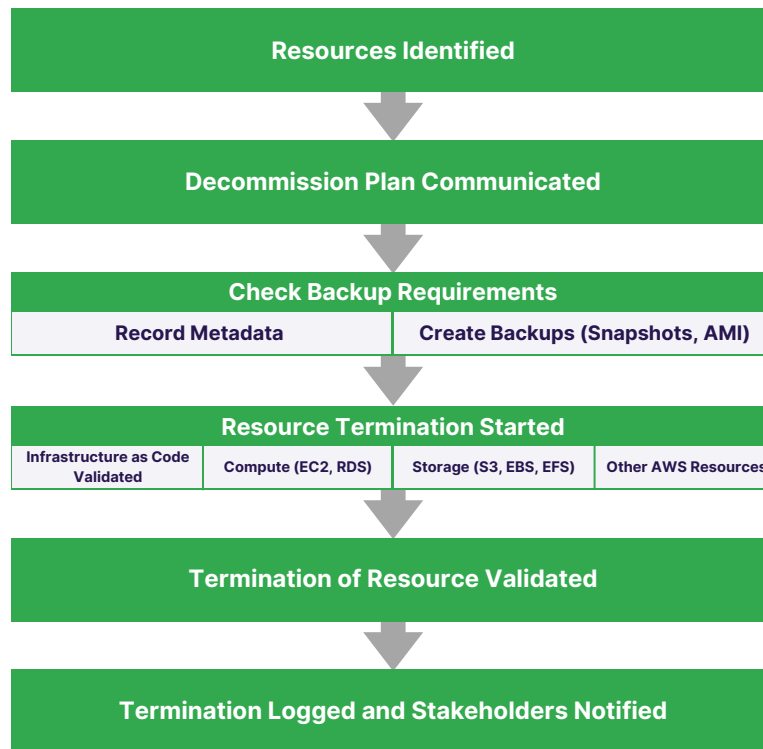
Don't let idle or unused resources become forgotten zombies in your cloud that drain your budget. Remember, the best time to decommission is before the bill arrives!

To decommission resources, you need to:

- **Track Resources Over Their Lifetime:** Implement resource tagging and monitoring to identify idle instances, underutilized databases, and forgotten projects.
- **Create a Decommissioning Playbook:** Establish a clear process for identifying and terminating unused resources. Before decommissioning, verify that the resources you have identified for decommissioning are not being utilized.



A flowchart displaying the steps in the decommissioning process.



 **Choosing Cost-Optimized Resources**

Cost savings relies on selecting suitable services, resources, and configurations for your workloads. If you want the biggest bang for your buck, then take into account the following factors when choosing AWS resources.

 **Consider Costs When Selecting Services**

When selecting services within AWS, it's imperative to evaluate costs effectively to ensure optimal resource allocation and budget management. Here's a breakdown of the process:

Identify Your Business's Cost Requirements

- Collaborate with your team to review each workload and define the balance between cost optimization and other pillars such as performance and reliability.
- You can address any technical debt identified and plan future optimization through documentation of backlog items to revisit.

Analyze All Components of the Workload

- Break down the workload into its individual components (i.e. web servers, application servers, databases, caches, queues, load balancers, etc.)

- Prioritize these components in order of effort as defined by your organization’s priorities.
- Review the AWS services available for each component and choose the option that aligns best with your business priorities.
- Evaluate the overall cost to the organization of each component and calculate the TCO.

 **Choose the Correct Resource Type, Size, Number, and the Best Pricing Model**

You can select the appropriate resource type, size, number, and pricing model to fulfil your technical requirements at the lowest cost. Implement the following steps:

Cost Modeling:

- Analyze the total cost of your entire workload and each individual component.
- Perform benchmarking tests under various predicted loads to compare costs and identify optimal configurations.

Resource Type:

Select resource type based on data about the workload and resource characteristics. To determine the needed instance type:

- Gather details about the system requirements of your application or software that you plan to run on your instance. These details should include the following:

Operating system	Number of CPU cores	GPU cores	Amount of system memory (RAM)	Storage type and space	Network bandwidth requirement
------------------	---------------------	-----------	-------------------------------	------------------------	-------------------------------

- Identify compute requirements and explore Amazon EC2 instance families such as:

General Purpose	Compute Optimized	Memory Optimized	Storage Optimized	Accelerated Computing	HPC Optimized
-----------------	-------------------	------------------	-------------------	-----------------------	---------------

Resource Size and Number:

- **Rightsizing:** Analyze workload metrics to avoid over-provisioning or underutilization. Choose the most appropriate instance size for optimal performance.
- **Auto-Scaling:** Implement auto-scaling policies to dynamically adjust resource allocation based on workload demand.

AWS Pricing Models:

You can leverage one or a combination of pricing models to minimize resource costs:

AWS Pricing Option	Description	Approximate Savings (%)	Use Cases
On-Demand Instances	<ul style="list-style-type: none"> • Default, pay as you go pricing model • Charged a flat hourly rate • No upfront commitments • Immediate access to resources 	Limited	<ul style="list-style-type: none"> • Applications with short-term workloads • Unpredictable workloads that can't be interrupted • Pre-production environments
Reserved Instances	<ul style="list-style-type: none"> • Reserved for a specific period (1 or 3 years) • Upfront or partial upfront payment 	Up to 72%	<ul style="list-style-type: none"> • Steady-state or predictable workloads • Production environments • Long-term projects
Savings Plans	<ul style="list-style-type: none"> • Flexible plan committing to a specified amount of compute usage (measured in dollars per hour) for 1 or 3 years • Do not provide capacity reservations 	Up to 72%	<ul style="list-style-type: none"> • Multi-service usage • Dynamic workloads • Variable workloads
Spot Instances	<ul style="list-style-type: none"> • Bid for and use when spot price is lower • Can be interrupted or terminated by AWS when spot price exceeds bid price or during high demand for on-demand instances 	Up to 90%	<ul style="list-style-type: none"> • Batch processing • Testing and development environments • Stateless workloads



Optimize Data Storage and Transfer Costs

Data storage and transfer often sneak up as significant cost contributors in the cloud. Answer these questions to segment your data and assess your storage needs:

- How frequently and rapidly do you access your data?
- Does your data storage demand high IOPS or throughput?

Next, consider these aspects:

<p>Data Importance ×</p> <p>Critical or regulated data may require long-term, high-cost storage.</p>	<p>Data Sensitivity ×</p> <p>Highly sensitive data needs robust protection against changes, loss, and corruption.</p>	<p>Data Size ×</p> <p>Knowing your data's size helps estimate storage needs and costs.</p>	<p>Data Lifespan ×</p> <p>Transient data is short-lived and doesn't demand high durability.</p>
---	--	---	--

Amazon's data transfer rates vary based on transfer types. Transfers between AWS Regions are the costliest, followed by transfers between Availability Zones (AZs). The lowest rates apply to transfers within a single AZ.

Here's your playbook to combat data storage and transfer costs:

- **Tier Strategically:** Utilize storage classes based on access frequency. Park rarely accessed data in cost-effective Glacier or Deep Glacier tiers.
- **Lifecycle Management:** Set automated rules to automatically move data between tiers based on access patterns.
- **Compress Data:** Reduce storage footprint with compression tools for backups, logs, and archival data.
- **Delete Unnecessary Data:** Regularly audit and remove redundant or obsolete data.
- **Transfer Within the Same Region: Avoid** inter-region data transfer costs by keeping data and compute resources in the same region.
- **Utilize Amazon CloudFront:** This content delivery network caches data closer to users, reducing data transfer costs and latency.
- **Batch Transfers:** Transfer large datasets infrequently instead of smaller, frequent transfers.
- **Compressed Data transfers:** Reduce transfer size by compressing data before transferring it.
- **Dedicated Connections:** Explore dedicated network connections for predictable, high-volume data transfers.



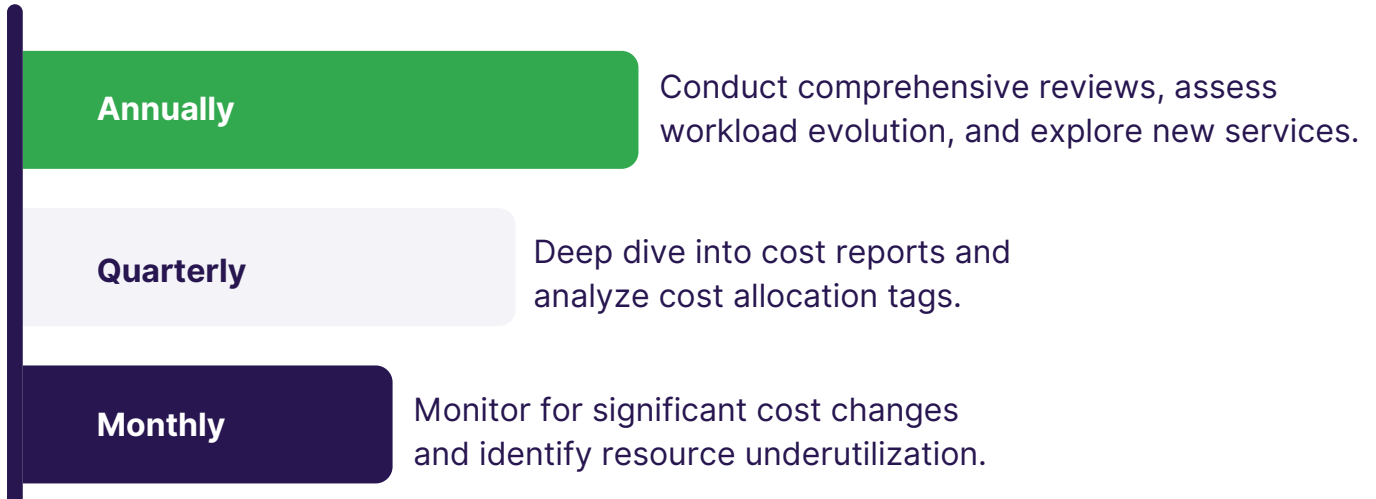
Continuously Optimizing Over Time

With each release of new services and features by AWS, it's crucial to revisit your existing architectural choices to ensure that they remain cost-effective. As your needs evolve, adopt a proactive approach to decommissioning resources that are no longer necessary. You can embrace the following best practices:

Develop a Workload Review Process

At this stage, we advise you to create a procedure outlining criteria and frequency for workload evaluation.

- **Schedule Recurring Reviews:**



- **Define Review Thoroughness:** Decide how much effort your team should spend on the review of the workload or workload components.

Perform Continuous Analysis on the Workload Demand

Workload demands evolve over time, and so should your cost optimization strategies. Continuously analyzing your workload helps to identify and address potential cost leaks before they become significant.

When analyzing workload demand for cost optimization, you can consider:

Resource usage and performance metrics	User and application scaling patterns	Types of workloads	Service-level agreements (SLAs)
--	---------------------------------------	--------------------	---------------------------------

Here's how you can stay ahead of the curve:

- **Historical Analysis:** Uncover workload usage patterns through Amazon CloudWatch logs and historical data. Also, consider data from seasonal trends and previous iterations.
- **Predict External Influences:** Collaborate with cross-functional (sales, marketing, or business development etc.) that can impact workload demand. Forecast the workload demand with this data.



Automate Operations for Efficiency

Introducing automation can significantly reduce the time and effort required for administrative operations. Automating tasks such as provisioning resources, managing deployments, and scaling infrastructures become more cost-effective and less prone to errors.

- **Deploy Once, Scale Endlessly:** Leverage Infrastructure as Code (IaC) tools such as AWS CloudFormation, AWS SDK, AWS CLI, or Terraform to enable seamless deployments for consistency and repeatability.
- **Streamline Operations:** Automate routine tasks without manual intervention. Customize services and tools according to your unique business requirements.

The Smart Way to Cloud: Your AWS Cost
Optimization Handbook

Taking the Next Step in Your Cost Optimization Journey

www.cloudelligent.com

Customer Case Study




Our customer is an analytics consulting firm that guides their clients to modernize data and analytics processes.

Cloudelligent experts regularly audits their AWS infrastructure to identify and optimize their resource usage. We rightsize resources according to customer needs or decommission them when they are no longer needed.

By implementing these proactive measures, the customer achieved significant cost savings of **26%**

Accelerate Your Journey To Smart Cloud Cost Management With Cloudelligent

Seize the opportunity to Spend Smarter, Secure Success, and Start Saving with Cloudelligent. Our expert team is here to guide your business through the complexities of the cloud costs. Take the first step toward AWS cost excellence with our FREE Cost Optimization Assessment.


Ready to Spend Smarter?
Get a FREE Cost Optimization Assessment!
