

# Automating AWS: Right-Sizing with AI and Machine Learning

Mohit Thodupunuri

**Abstract:** *As cloud adoption accelerates, organizations face the challenge of balancing performance with cost efficiency in AWS environments. Right-sizing—selecting the optimal computing, storage, and networking resources—plays a critical role in minimizing waste while ensuring workload performance. However, manual right-sizing is complex, time-consuming, and prone to inaccuracies. This paper explores how Artificial Intelligence (AI) and Machine Learning (ML) can automate AWS right-sizing, improving efficiency, reducing costs, and enhancing scalability.*

**Keywords:** AWS cost optimization, AI-based right-sizing, machine learning cloud management, automated AWS resources, intelligent workload optimization

## 1. Introduction

Cloud computing has transformed how businesses manage their IT resources. Amazon Web Services (AWS), a leader in this space, offers incredible flexibility and scalability. However, with flexibility comes complexity—especially in managing costs effectively.

One of the most crucial yet overlooked cost-saving strategies in AWS is right-sizing. Right-sizing ensures that cloud resources match actual usage, avoiding both under-provisioning and over-provisioning. But, doing this is a time-consuming and error-prone task. That's where automation comes into play. More specifically, automation is powered by Artificial Intelligence (AI) and Machine Learning (ML). These technologies are revolutionizing how organizations optimize cloud usage. By using intelligent models, businesses can analyze usage patterns and predict future needs. Consequently, they can make informed decisions without manual guesswork. This approach not only saves money but also boosts performance and operational efficiency.

Traditionally, right-sizing required constant monitoring of resource consumption. Cloud engineers would analyze CPU utilization, memory usage, and disk I/O. Then, they would match this data to the best-fitting instance type. Sounds simple, right? In reality, it's far from it. Human oversight, fluctuating workloads, and fast-changing environments make this task difficult to scale. Additionally, with hundreds or even thousands of instances running, manual right-sizing quickly becomes unmanageable.

AI and ML change this equation entirely. These technologies learn from historical data and recognize usage trends over time. As a result, they can recommend optimal configurations for each workload. Moreover, they continuously adapt as demand shifts. This means businesses no longer have to rely on fixed schedules or static thresholds. Instead, right-sizing becomes a dynamic, intelligent process.

Furthermore, integrating AI and ML into right-sizing tools provides predictive capabilities. Rather than reacting to usage spikes or underutilization, the system can proactively adjust resources. For instance, if a model predicts lower usage during weekends, it can automatically scale down non-critical resources. On the other hand, it can scale up in anticipation of known traffic increases. This level of automation brings significant improvements in both cost control and performance optimization.

Let's also not forget about visibility and transparency. Machine learning models can offer detailed reports and insights. These reports help teams understand not only what's changing but also why it's changing. Such clarity is essential for building trust in automated systems. In turn, this fosters more widespread adoption across departments.

Implementing AI-driven right-sizing requires upfront investment in the form of quality data, robust algorithms, and integration with existing AWS services. However, the long-term gains far outweigh the initial effort. Companies that embrace AI and ML in cloud cost optimization see double-digit savings within months. And with AWS offering native tools like Compute Optimizer and Trusted Advisor, the barrier to entry is lower than ever.

The combination of automation, AI, and machine learning is a game-changer for AWS right-sizing. It transforms a complex, ongoing task into a streamlined, intelligent process.

## 2. Literature Review

Right-sizing in cloud environments, particularly within Amazon Web Services (AWS), is a critical practice aimed at aligning computing resources with actual workload demands. Traditional approaches to right-sizing rely heavily on manual monitoring and static rules. However, these methods often fall short when applied to dynamic and large-scale cloud environments. As a result, the integration of Artificial Intelligence (AI) and Machine Learning (ML) has emerged as a powerful solution for automating this process.

Recent studies demonstrate that ML-based models significantly enhance resource allocation by predicting future usage patterns and adjusting resources accordingly [1]. These models rely on historical usage data to learn behavior over time, enabling more accurate and timely adjustments than rule-based systems. In turn, this leads to improved efficiency and performance across cloud infrastructures [2]. The implementation of intelligent automation helps eliminate the guesswork traditionally associated with right-sizing while reducing the likelihood of overprovisioning or resource shortages.

Furthermore, the development of tools such as AWS Compute Optimizer has made it easier for organizations to adopt automated right-sizing strategies. These tools leverage ML to analyze CPU, memory, and network utilization, subsequently recommending optimal instance types and configurations [3][9][10]. The automation of these recommendations reduces operational overhead and simplifies decision-making for cloud architects and DevOps teams.

In addition to these built-in services, newer platforms are advancing the concept by enabling real-time AI-powered decision-making. These platforms observe application behavior continuously and apply right-sizing recommendations proactively rather than reactively [4]. This approach offers greater flexibility and responsiveness to workload fluctuations, enhancing both performance and cost efficiency.

Right-sizing frameworks are also increasingly guided by comprehensive best practices that emphasize continuous monitoring, usage forecasting, and adaptive scaling [5][6]. The literature suggests that integrating AI allows these best practices to evolve into self-optimizing systems that respond dynamically to changes in resource demand. In contrast to static threshold-based triggers, AI-powered solutions can make nuanced decisions based on multi-dimensional data analysis.

Moreover, advanced AI techniques, including generative models, are being explored to simulate potential usage scenarios before they occur. This proactive simulation enables more informed resource planning and preemptive right-sizing, especially in environments with seasonal or unpredictable demand patterns [7]. By anticipating future needs, these models help organizations maintain cost control without compromising application performance.

It is also important to note that effective right-sizing is not solely about minimizing costs. Maintaining service quality and supporting business continuity remain essential considerations. Studies indicate that AI-enhanced automation achieves this balance more effectively than manual methods, as it factors in both technical performance and business requirements [8].

Taken together, the reviewed literature highlights a growing consensus: the automation of right-sizing using AI and ML offers substantial benefits over traditional methods. These technologies provide scalable, intelligent, and cost-effective solutions that support the evolving demands of modern cloud computing environments.

### 3. Problem Statement: The Inefficiency of Manual AWS Resource Allocation

#### 3.1 Complexity of AWS Resource Selection:

One of the most pressing challenges in manual resource allocation is navigating the complexity of AWS offerings. AWS provides hundreds of instance types, various storage classes, and intricate networking configurations. While this diversity allows for tailored infrastructure solutions, it simultaneously creates a burdensome decision-making process.

Manually evaluating these configurations demands a high level of expertise and a deep understanding of each application's behavior. Moreover, staying current with AWS's frequent updates and new service offerings requires continuous learning and adaptation. This complexity makes it difficult for teams to consistently select the optimal combination of resources, especially when managing large-scale or multi-region deployments.

#### 3.2 Time-Consuming Manual Analysis

Another significant drawback of manual AWS resource allocation is the time investment required. Teams must regularly collect and review performance metrics from sources like AWS CloudWatch, Trusted Advisor, and third-party monitoring tools. This process often involves sifting through large volumes of data to identify patterns in CPU usage, memory consumption, and I/O operations.

Even more time-consuming is the task of correlating these metrics with actual resource utilization to identify inefficiencies or bottlenecks. As a result, valuable operational hours are consumed by repetitive, labor-intensive analysis—time that could otherwise be spent on innovation or strategic improvements.

#### 3.3 Inaccuracy and Human Error

Manual approaches are prone to human error and subjective decision-making. Despite best intentions, analysts may misinterpret performance metrics or apply generalized assumptions to specific workloads. This can lead to overprovisioning—where more resources than necessary are allocated—resulting in inflated operational costs.

Conversely, under-provisioning can restrict application performance, causing slow response times, degraded user experiences, and even system downtime. These inaccuracies are amplified in fast-growing environments where decisions must be made quickly and frequently. In such scenarios, reliance on manual methods introduces unnecessary risk and inefficiency.

#### 3.4 Dynamic Workload Variability

Manual resource allocation is inherently unsuited for dealing with the dynamic nature of cloud workloads. Demand for cloud applications fluctuates based on time of day, user behavior,

seasonal trends, and other unpredictable factors. Static provisioning models cannot accommodate these variations effectively.

Manual adjustments are reactive and often delayed—leaving systems either underprepared for peak loads or wasting resources during off-peak periods. Without the ability to anticipate and respond to real-time changes, organizations face performance bottlenecks, service interruptions, or escalating cloud bills.

Manual AWS resource allocation is hindered by the sheer complexity of options, the time-intensive nature of analysis, the risk of human error, and the inability to adapt quickly to changing workloads. These limitations underscore the urgent need for automated solutions driven by AI and machine learning that can streamline right-sizing decisions and improve operational efficiency.

#### 4. Solution: AI and ML-driven automation for AWS Right-sizing

In today's rapidly evolving cloud landscape, businesses require more than just access to powerful infrastructure—they need intelligent, adaptive systems that can scale and optimize resources in real time. AWS offers a flexible platform, but managing it manually leads to inefficiencies, cost overruns, and performance challenges. By integrating Artificial Intelligence (AI) and Machine Learning (ML) into AWS right-sizing, organizations can eliminate guesswork, accelerate decision-making, and dramatically reduce cloud waste. These technologies enable systems to automatically analyze historical and real-time data, identify optimal configurations, and adapt continuously to workload fluctuations. The result is a more responsive, cost-effective, and performance-oriented cloud environment.

##### 4.1 Predictive Modeling for Resource Demand

Predictive modeling lies at the heart of intelligent right-sizing. ML algorithms can be trained on historical resource usage data, learning the behavior of workloads over time. These models detect trends and seasonality and then forecast future usage with impressive accuracy. As a result, systems can scale proactively, allocating resources before demand peaks. This foresight reduces performance bottlenecks and avoids sudden spikes in cost due to reactive scaling.

Consider the following example using Amazon SageMaker and XGBoost to predict CPU utilization based on CloudWatch metrics:

```
import pandas as pd

from sagemaker import Session

from sagemaker.xgboost import XGBoost

from sagemaker.inputs import TrainingInput

# Assume 'df' is a DataFrame with CloudWatch data
df.to_csv('cpu_data.csv', index=False)

session = Session()

bucket = session.default_bucket()

prefix = 'aws-right-sizing-model'

s3_input_path = f's3://{bucket}/{prefix}/cpu_data.csv'

session.upload_data('cpu_data.csv', bucket=bucket, key_prefix=prefix)

xgb = XGBoost(entry_point='train.py', framework_version='1.3-1', role='SageMakerRole',
              instance_count=1, instance_type='ml.m5.large')

xgb.fit({'train': TrainingInput(s3_input_path, content_type='csv')})
```

**Figure 1:** Using Amazon SageMaker and XGBoost to predict based on CloudWatch metrics

##### 4.2 Automated Resource Recommendation Engines

AI-powered engines go a step further by providing actionable resource recommendations. These systems monitor metrics such as CPU utilization, memory usage, disk I/O, and network throughput in real-time. They then suggest optimal EC2 instance types, EBS volume configurations, or load balancer adjustments. Unlike manual methods, AI doesn't rely on static rules—it continuously learns and improves.

AWS Compute Optimizer is an example of this approach. It evaluates your current deployments and compares them to usage patterns to recommend more efficient configurations. Here's a sample Python script to fetch recommendations using Boto3:

```
import boto3

client = boto3.client('compute-optimizer')

response = client.get_ec2_instance_recommendations(
    accountIds=['123456789012'],
    filters=[{'name': 'Finding', 'values': ['Overprovisioned']}],
    maxResults=10
)

for recommendation in response['instanceRecommendations']:
    print(f"Instance ID: {recommendation['instanceArn']}")
    for option in recommendation['recommendationOptions']:
        print(f"Suggested Instance Type: {option['instanceType']},
              Projected Savings: {option['performanceRisk']}")
```

**Figure 2:** A sample Python script to fetch recommendations using Boto3

### 4.3 Continuous Optimization and Feedback Loops

AI and ML systems thrive when they are allowed to evolve. Through continuous optimization, they refine predictions and recommendations based on real-world outcomes. Feedback loops provide critical insights—did the last recommendation save money? Did the new instance perform better or worse? The system incorporates this feedback and retrains models, improving future decisions.

A Lambda function can automate this feedback loop. After every resource change, it logs performance and cost data. That data is then used to retrain the ML model at regular intervals.

```
import boto3

def lambda_handler(event, context):
    cloudwatch = boto3.client('cloudwatch')
    metrics = cloudwatch.get_metric_statistics(
        Namespace='AWS/EC2',
        MetricName='CPUUtilization',
        Dimensions=[{'Name': 'InstanceId', 'Value': 'i-0123456789abcdef0'}],
        StartTime='2025-04-01T00:00:00Z',
        EndTime='2025-04-10T00:00:00Z',
        Period=3600,
        Statistics=['Average']
    )
    # Save metrics for retraining ML model
    save_metrics_to_s3(metrics)
```

**Figure 3:** Retraining ML models using a Lambda function

Over time, the system becomes smarter, more accurate, and highly efficient at managing AWS resources.

### 4.4 Automated Scaling and Deployment

One of the most transformative applications of AI and ML is the ability to act autonomously. Automated scaling ensures the environment adjusts to demand instantly—no waiting for a human to respond. ML models assess workload pressure and trigger auto-scaling groups to spin up or scale down resources as needed. Combined with Infrastructure as Code (IaC), this allows dynamic deployment of the right resources in the right place at the right time.

Using AWS Auto Scaling and ML integration, a system can scale an EC2 fleet based on CPU predictions. A pre-trained ML model outputs expected demand. A Lambda function then updates the Auto Scaling policy accordingly.

```
import boto3

autoscaling = boto3.client('autoscaling')

autoscaling.update_auto_scaling_group(
    AutoScalingGroupName='web-app-asg',
    MinSize=2,
    MaxSize=10,
    DesiredCapacity=predicted_capacity
)
```

**Figure 4:** Updating Auto Scaling policy with a Lambda function

This approach eliminates the lag between resource need and availability, ensuring consistent performance and cost control.

Automating AWS right-sizing with AI and machine learning transforms a reactive, manual process into a proactive, intelligent system. It not only boosts performance and lowers costs but also frees up teams to focus on innovation rather than maintenance. Through automated recommendations, adaptive learning, and dynamic scaling, organizations can unlock the full potential of their AWS infrastructure.

## 5. Recommendation: Implementing and Optimizing AI-Driven Right-sizing Strategies

With businesses transitioning to increasingly complex cloud infrastructures, the need for precise and efficient resource allocation becomes even more critical. Manual approaches are no longer scalable or sustainable in environments that evolve by the second. Artificial Intelligence (AI) and Machine Learning (ML) offer a powerful solution to these challenges.

However, the successful implementation of AI and machine learning requires strategic planning, continuous optimization, and thoughtful integration with AWS services. Organizations must focus not only on deploying AI models but also on aligning them with real-world use cases, infrastructure constraints, and business goals. This section outlines key recommendations to help teams effectively implement AI-driven right-sizing for AWS.

### 5.1 Selecting Appropriate ML Algorithms

Choosing the right ML algorithms is foundational to the success of any AI-driven right-sizing strategy. Different workloads exhibit different patterns, and not all algorithms are suitable for all scenarios. For example, workloads that show seasonal or time-based spikes in demand benefit from time-series forecasting models like ARIMA or Facebook Prophet. These models learn from historical usage trends to make forward-looking predictions about resource needs.

In contrast, workloads with relatively static usage patterns may benefit more from regression models that predict performance metrics based on existing configurations. Clustering algorithms like K-means can be extremely useful for grouping similar



usage patterns. These models allow the system to recommend configurations based on observed behaviors from similar applications. Selecting algorithms that align with the workload profile ensures higher accuracy and better performance outcomes. Therefore, investing time in algorithm evaluation and experimentation is well worth the effort

### 5.2 Data Collection and Preprocessing

High-quality predictions begin with high-quality data. AI and ML systems rely heavily on accurate, consistent, and comprehensive datasets. Raw data collected from AWS services like CloudWatch often includes noise, missing values, and inconsistencies. If left unaddressed, these issues can severely compromise model performance. Data preprocessing plays a crucial role here.

Normalization, outlier detection, and missing value imputation are among the most common preprocessing techniques used to clean and prepare data for training. Feature engineering, such as creating composite metrics from CPU and memory usage, can also improve model accuracy. Establishing a robust data pipeline that automatically handles data cleaning ensures that the models always learn from the best available inputs. Ultimately, the quality of data directly influences the reliability of the right-sizing recommendations.

### 5.3 Integration with AWS Services

Implementing AI-driven right-sizing strategies in isolation is not enough. The solution must integrate seamlessly with AWS services to deliver real-world impact. CloudWatch provides the performance metrics that feed predictive models, while Auto Scaling handles the provisioning and de-provisioning of resources based on those predictions. AWS Lambda can automate the execution of these workflows, allowing real-time adjustments without human intervention.

By integrating the ML system with AWS Step Functions, organizations can orchestrate complex workflows that include monitoring, prediction, scaling, and logging. This interconnected approach ensures that right-sizing happens continuously and automatically. It eliminates manual intervention and reduces the risk of oversight. Furthermore, leveraging services like AWS SageMaker for model deployment brings scalability and reliability to the ML lifecycle.

### 5.4 Continuous Monitoring and Evaluation

Deploying a model is not the end of the journey—it's just the beginning. Continuous monitoring is essential to ensure that the AI-driven right-sizing system remains effective over time. Cloud environments are dynamic, and workload behaviors evolve. A model that works well today may underperform tomorrow if not regularly evaluated and updated.

Monitoring tools must track both technical performance metrics and business outcomes. This includes evaluating how often

recommendations are accepted, measuring cost savings achieved, and monitoring system performance before and after right-sizing actions. Feedback loops should be used to retrain the models with new data, adapting to changes in workload patterns. Regular audits and performance reviews help refine the system and identify areas for improvement. Without continuous evaluation, even the most sophisticated AI solutions can drift and lose relevance.

## 6. Conclusion

Right-sizing in AWS is no longer a task that can be effectively managed with manual effort alone. The scale, speed, and complexity of modern cloud infrastructures demand a smarter, more adaptive approach. AI and Machine Learning offer that edge by turning historical and real-time data into actionable insights. However, realizing the full potential of AI-driven right-sizing requires a holistic strategy—selecting the right algorithms, preparing the right data, integrating with AWS tools, and continuously monitoring results.

Organizations that invest in these strategies position themselves for long-term cloud success. They minimize waste, maximize performance, and gain the agility needed to innovate faster. As AWS environments continue to grow, automation through AI will not just be a recommendation—it will be a necessity.

## References

- [1] Andrea Cesarini, "Optimizing Cloud Resource Allocation Using Machine Learning Techniques", *Perspective Journal of Computer Science & Systems Biology*, Vol 17, p 512, 2024.
- [2] Sadia Syed and Eid Mohammad Albalawi, "Optimizing Cloud Resource Allocation with Machine Learning: A Comprehensive Approach to Efficiency and Performance", *ResearchGate Preprint*, DOI: 10.21203/rs.3.rs-4825637/v1, August 2024.
- [3] Cody Slingerland, "What Is AWS Compute Optimizer? A Newbie-Friendly Guide", *CloudZero Blog*, March 20, 2024.
- [4] John Jamie, "Introducing AI-Powered Right-sizing for AWS EC2 VMs", *Sedai Blog*, May 7, 2024.
- [5] nOps, "The Definitive Guide to AWS Right-sizing", *nOps eBook*, 2024.
- [6] Cody Slingerland, "What Is AWS Sizing? 10+ Best Practices And Tips", *CloudZero Blog*, June 12, 2024.
- [7] Matthias Patzak, "Generative AI cost optimization strategies", *AWS Enterprise Strategy Blog*, September 23, 2024.
- [8] Jana Brnakova, "Right-size your AWS infrastructure: optimize costs without compromising your business", *Revolvy Blog*, January 11, 2024.
- [9] AWS Documentation, "What is AWS Compute Optimizer?", *AWS Compute Optimizer User Guide*, 2024.
- [10] AWS, "AWS Compute Optimizer", *AWS Product Page*, 2024.