

Unlocking the Power of Cloud Commitment Discounts: A Deep Dive into Reserved Instances and Savings Plans

Author: Venkata Sasidhar(Sasi) Kanumuri,
Company: Sunnyslabs AI Corp

Abstract

Moving data and applications from traditional on-premises data centers to cloud infrastructure offers companies the potential for significant cost savings through accelerating innovation, keeping a competitive edge, and better interacting with customers and employees. You can scale your cloud resources up or down to meet demand, and costs will follow. However, cloud service costs can be higher than anticipated, so monitoring and optimizing your cloud spend is critical. Cloud cost optimization combines

Keywords: reserved instances, savings plan, commitment discounts, sustained usage discounts, AWS SP, RI, GCP SP, Azure SP, Azure Reservations, storage reservations, cloud spending, cloud cost optimization

INTRODUCTION

Navigating the ever-shifting landscape of cloud costs can be a daunting task. While the cloud offers undeniable scalability and agility, keeping expenses in check requires active management and strategic planning. This article dives deep into two powerful tools for cloud cost optimization: Reserved Instances (RIs) and Savings Plans (SPs). We'll explore how these commitment-based offerings can unlock significant discounts on compute costs, compare their strengths and limitations across major cloud providers like AWS, Azure, and GCP, and equip you with additional strategies for maximizing your cloud budget. Whether you're an IT professional managing complex cloud environments or a user seeking cost-effective solutions, this guide will empower you to make informed decisions and navigate the path toward optimal cloud spending.

strategies, techniques, best practices, and tools to help reduce cloud costs, find the most cost-effective way to run your applications in the cloud environment, and maximize business value. Monitoring metrics and comparing data when using multiple cloud vendors with different dashboards can be difficult, and overspending can be easy. It's about reducing costs and ensuring they align with your business goals. In other words, paying more may make sense if you earn more revenue or see more productive activities and profitability from a cloud service provider.

1. Understanding Savings Plans vs. Reserved Instances

A. Savings Plan

A Savings Plan is another flexible pricing model that provides up to 72% savings on your AWS compute usage. Savings Plans offer significant savings over On-Demand Instances, just like EC2 Reserved Instances, in exchange for a commitment to use a specific amount of compute power (measured in \$/hour) for one or three years. You can sign up for Savings Plans for a one- or three-year term and easily manage your plans by taking advantage of recommendations, performance reporting, and budget alerts in the AWS Cost Explorer. Note - Savings Plans do not provide a capacity reservation. You can, however, reserve capacity with On Demand Capacity Reservations and pay lower prices on them with Savings Plans. You can continue purchasing RIs to maintain compatibility with your existing cost management processes, and your RIs will work alongside Savings Plans to reduce your overall bill. However, as your RIs expire, we encourage you to sign up for Savings Plans, which offer the same savings as RIs but with additional flexibility.

AWS offers three types of Savings Plans:

- **Compute Savings Plans** provide the most flexibility and help to reduce your costs by up to 66% (just like Convertible RIs). These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, Region, operating system, or tenancy, and

also apply to Fargate and Lambda usage. For example, with Compute Savings Plans, you can change from C4 to M5 instances, shift a workload from EU (Ireland) to Europe (London), or move a workload from Amazon EC2 to Fargate or Lambda at any time and automatically continue to pay the Savings Plans price.

- EC2 Instance Savings Plans provide the lowest prices, offering savings of up to 72% (just like Standard RIs) in exchange for a commitment to the usage of individual instance families in a Region (for example, M5 usage in N. Virginia). This automatically reduces your cost on the selected instance family within that region, regardless of instance size, operating system, or tenancy (Shared, Dedicated, or Outpost). This flexibility allows you to optimize your computing environment within the chosen family without sacrificing cost savings. For instance, you could switch from an m5.xlarge Linux instance to a c5.2xlarge Windows instance and still benefit from the Savings Plans discount, as long as both instances belong to the same family (e.g., M or C).
- Amazon Sage Maker Saving Plans: Machine learning workloads running on SageMaker are catered explicitly to by Amazon SageMaker Savings Plans. They provide savings of up to 64% in return for a one- or three-year commitment to a regular hourly use level, much to Compute Savings Plans. The SageMaker Studio Notebooks, SageMaker On-Demand Notebooks, SageMaker Processing, SageMaker Data Wrangler, SageMaker Training, SageMaker Real-Time Inference, and SageMaker Batch Transform are among the qualifying SageMaker ML instances to which these plans automatically apply. This flexibility within SageMaker helps you to optimize expenses throughout your machine learning development and deployment pipeline. Savings Plans is another flexible pricing model that provides up to 72% savings on your AWS compute usage. This pricing model offers lower prices on Amazon EC2 instances usage, regardless of instance family, size, OS, tenancy, or AWS Region, and also applies to AWS Fargate and AWS Lambda usage. Although AWS provides Savings Plans, comparable alternatives for optimizing computation costs are available from other major cloud providers, such as Microsoft Azure and Google Cloud Platform (GCP). These services from GCP and Azure will be covered later.

B. Reserved Instances

A Reserved Instance is not an instance dedicated to your organization. It is a billing discount applied to using On-Demand Instances in your account. These On-Demand Instances must match certain attributes of the Reserved Instances you purchased to benefit from the billing discount. You pay for the entire term of a Reserved Instance, regardless of actual usage,

so your cost savings are closely tied to use. Therefore, planning and monitoring your usage is vital to maximize your investment. When you purchase Reserved Instances, you make a one-year or three-year commitment and receive a billing discount of up to 72 percent in return. Reserved Instances can save you a lot of money when used for the appropriate workloads. When you purchase a Reserved Instance in a specific Availability Zone, it provides a capacity reservation. This improves the likelihood that your computing capacity will be available in a particular Availability Zone when needed. A Reserved Instance purchased for an AWS Region does not provide capacity reservation. Limitations, for instance, size flexibility, do not apply to the following Reserved Instances:

- Reserved Instances that are purchased for a specific Availability Zone (zonal Reserved Instances)
- Reserved Instances with dedicated tenancy
- Reserved Instances for Windows Server, Windows Server with SQL Standard, Windows Server With SQL Server Enterprise, Windows Server with SQL Server Web, RHEL, and SUSE Linux Enterprise Server
- Reserved Instances for G4ad, G4dn, G5, and G5g instances.

Cloud providers offer various options for optimizing compute costs. While Amazon Web Services (AWS) utilizes Reserved Instances with significant discounts, Google Cloud Platform (GCP) takes a different approach. GCP offers Committed Use Contracts (CUCs) that achieve similar cost savings. Through CUCs, you can reserve virtual machines (VMs) for a set period (usually 1 or 3 years) and receive discounted pricing. This approach ensures guaranteed VM availability in your chosen region and zone, alongside cost reductions compared to on-demand pricing.

Here are some key benefits of GCP's Committed Use Contracts:

- Discounted Pricing: Committing to a CUC term allows you to benefit from significant discounts compared to regular on-demand pricing.
- Instance Reservations: CUCs ensure reserved VMs are always available within your specified region and zone, meeting your resource needs.
- Usage Flexibility: Despite reserving VMs, CUCs offer flexibility when used across different projects within your GCP account.

GCP's Committed Use Contracts might evolve. For the latest information and pricing details, refer to the official Google Cloud documentation or consult their sales team.

With Reserved Instances, you can save up to 72% over equivalent on-demand capacity. Reserved Instances are available in 3 options – All up-front (AURI), partial up-front (PURI), or no upfront payments (NURI). When you buy Reserved Instances, the larger the upfront payment, the greater the discount. To maximize your savings, you can

pay everything up-front and receive the largest discount. Partial up-front RIs offer lower discounts but allow you to spend less upfront. Lastly, you can spend nothing upfront and receive a smaller discount, freeing up capital to spend on other projects.

2. Types of GCP Discounts:

Users can leverage multiple types of GCP discounts. We'll explain each one in detail in these sections.

A. Sustained-use Discounts

Compute Engine instances used more than 25% of the month will incur a sustained-use discount on

vCPU and memory costs. The longer the instance runs, the deeper the discount increases incrementally. It automatically applies to VMs that fall under Compute Engine and Kubernetes Engine. However, they do not apply to E2, A2, Tau T2D, and Tau T2A machine types. Discounts reset every first day of the month. Creating VMs on the first day of the month is advisable to maximize savings.

a. AWS versus GCP discount plans

Both AWS and Google Cloud offer discounts for one-year and three-year commitments, with few similarities beyond that.

You can see how they compare:

| Feature | AWS Standard RI | AWS Convertible RI | AWS Compute Savings Plan | AWS EC2 Instance Savings Plan | GCP CUD |
|-----------------------------------|---|--|---|-------------------------------------|---|
| Term | 1 or 3 years | 1 or 3 years | 1 or 3 years | 1 or 3 years | 1 or 3 years |
| Max discount | 72% | 66% | 66% | 72% | 70% |
| Payment options | All, partial or nothing upfront | As before | As before | As before | No upfront |
| What can you change? | AZ, instance type | Instance size and family, AZ, network type, OS, tenant | Instance size and family, AZ, network type, OS, tenancy, region | Instance size, AZ, network type, OS | Committing to a number of vCPUs and memory (not specific machine types) allows you to upgrade to new generation instances as they become available. |
| Further discount potential | No - remaining instances are charged at On-Demand rate. | As before | As before | As before | Sustained Use Discounts can be applied to any remaining instances. |

Fig 1.1 Table comparing AWS and GCP compute discount plans

b. Committed Use Discounts (CUDs)

Google Cloud answers the cloud savings question with committed use discounts (CUDs). Targeted at workloads with predictable resource needs, GCP CUDs allow you to purchase compute engine resources such as vCPUs, memory, GPUs, local SSDs, and sole-tenant nodes at a discounted rate if you commit to paying for those resources for either one or three years.

The discount ranges from up to 57% for resources, including machine types and GPUs, and up to 70% for memory-optimized machine types. You are billed monthly for the resources you bought for the term you chose (one or three years), even if you don't use those services, making them suitable for predictable workloads.

By analyzing your VM spending trends with and without a CUD, Google allows you to see how much you could save each month with a CUD contract. CUDs are available for single projects, or you can purchase multiple contracts and share them across numerous projects by enabling shared discounts.

CUDs are either spend-based or resource-based. Spend-based CUDs apply a discounted rate on usage based on a committed minimum dollar amount per

hour. Any usage over the committed amount is charged at the on-demand rate. Spend-based CUDs can be bought from your Cloud Billing account; they apply to eligible usage in any projects paid for by that Cloud Billing account. They are restricted to the following Google Cloud services like Cloud SQL, Google Cloud VMware Engine, Cloud Run, and Google Kubernetes Engine

Resource-based CUDs apply a reduced rate based on your commitment to use a minimum level of Compute Engine resources in a specific region. They can be used for vCPUs, memory, GPUs, and local SSDs. This allows you to use different machine instances across your organization when you have predictable Compute Engine resource needs.

B. Azure Reservations

Azure reservations are similar to savings plans and typically offer a more significant discount; however, they are limited to an extent. For virtual machines, the reservations must be used within the same instance size flexibility group or capacity priority (single scope only), and in the same region. In a nutshell, if your workload usage is going to be consistent (ie, stable) and no real changes to the

infrastructure are anticipated, then reservations are a good pick.

As with the savings plans, reservations can also be paid for monthly or upfront. Technically, reservations are also hourly, although the assumption is 730 hours/month of uptime. So what's meant by hourly, in this case, is if you were to put a VM that had a reservation applied to it into a stopped (deallocated) state, then the reservation would be released, and another VM in the same instance size flexibility group (if it is enabled on the reservation) would then make use of it IF one exists; if nothing else can utilize the reservation, then any hours it is not used will become lost (ie; unused reservation). If said VM is in a stopped state, it will still consume the reservation as compute hours are still charged in that situation. It's important to note that not all resources can have reservations applied to them. However, most reservations can be split, returned, or exchanged if they (or the finances) are needed elsewhere.

C. Azure Savings Plans

Azure savings plan is a method of obtaining resource cost reductions based on an hourly commitment charge. That commitment charge automatically applies to any compute resource within the savings plan's scope. It could apply to any compute resource if available and will not be locked down to a specific compute resource type (i.e., it won't just apply to VMs or app services).

It is important to note that not all resources can apply to savings plans, only compute resources as of this writing. If you have both reservations and a savings plan, the reservation will take precedence. If you have multiple savings plans, then the savings plan that offers the most savings will be used automatically (which is always the case with savings plans for computing). Anticipate your infrastructure or workloads changing relatively frequently, whether that means changing regions, VMs (outside of the flexibility group), etc. Savings plans will make the most sense in such situations.

Savings plans can be paid for monthly or upfront and can be applied to the following resources (as of this writing):

- Virtual Machines (only the compute costs, i.e., does not include storage, network, licensing, etc.)
- Azure Dedicated Host (as with VMs, only the compute costs are covered by the savings plan)
- Container Instances
- Azure Premium Functions
- Azure App Services (not all are included; Pv3 or isolated v2 is required)

NOTE:

Savings plans are not restricted to VM/instance series, OS, or regions, making them more flexible than reservations.

3. Comparison sheet across Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)

A comparison sheet focusing on reservations across Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), considering the following aspects:

- Services Covered Under Reservations: Lists the primary services for reservations on each cloud platform.
- Flexibility of Reservations: Describes the flexibility regarding reservation duration, instance types, and payment options.
- Most Discounted Instances Families: Highlights families or types that typically receive the highest discounts under reservation options.
- Estimated Savings for RIs: Provides a range of estimated savings percentages for Reservations.
- Estimated Savings for SPs: Provides a range of savings percentages for Savings Plans.
- Conversion Possibility: Indicates whether there are options to convert or modify existing reservations.
- Reselling RI/SP: Limited resale through select marketplaces. Check for restrictions and fees applied by both platforms and providers.

| Aspect | AWS | Azure | GCP |
|--|--|--|---|
| Services Covered Under Reservations | <p>Compute: Amazon EC2 Instances</p> <p>Databases: Amazon RDS, Amazon Redshift, Amazon DynamoDB</p> <p>Storage: Amazon EBS Volumes</p> <p>Other: Amazon ElastiCache, Amazon S3 Glacier</p> | <p>Compute: Azure Virtual Machines</p> <p>Databases: Azure SQL Database, Azure Cosmos DB</p> <p>Storage: Azure Disks, Azure Blob Storage</p> <p>Other: Azure App Service, Azure Kubernetes Service</p> | <p>Compute: Google Compute Engine</p> <p>Databases: Cloud SQL, Cloud Spanner</p> <p>Storage: Cloud Storage</p> <p>Other: Cloud Functions, Kubernetes Engine</p> |
| Flexibility of Reservations | Offers 1-year and 3-year terms with various payment options. Convertible reservations allow changing instance types. | Offers 1-year and 3-year terms with options for upfront, monthly, or no upfront payment. Azure Hybrid Benefit allows using on-premises licenses. | Offers 1-year and 3-year terms with various payment options. Sustained use discounts provide automatic savings. |
| Most Discounted Instances Families | Depending on the instance type and region, T3, M5, R5, etc. | D-series, E-series, F-series, etc., depending on the instance type and region. | N1, N2, E2, etc., depending on the instance type and region. |
| Estimated Savings for RIs | Up to 72%, with the highest discounts typically offered for sustained, predictable workloads on general-purpose (M, A) or compute-optimized (C) instance families. | Up to 73%, with Azure Reserved VM Instances with upfront payment offering the highest savings. Savings depend on instance type, region, term length, and upfront payment commitment. | Up to 60%, with sustained use discounts, typically offering the most significant savings for predictable workloads. |
| Estimated Savings for SPs | Compute Savings Plans: Up to 66%. | Reserved Virtual Machine Instances (RIs): Up to 72%. Azure Reserved VM Instances with upfront payment can offer the highest savings. | Committed Use Discounts (CUDs): Up to 50%. |
| Conversion Possibility | Convertible reservations allow changing instance types or platforms but may incur additional charges. | No direct conversion option, but reservations can be modified or exchanged within the same family. Azure Hybrid Benefit allows converting on-premises licenses. | No direct conversion option, but reservations can be modified or exchanged within the same family. Sustained use discounts are automatically applied. |

Fig. 1.2 - Table Comparison Among AWS, Azure, and GCP

This comparison provides a snapshot of the reservation options offered by AWS, Azure, and GCP, including their coverage, flexibility, discounts, and conversion possibilities. Businesses should assess their

needs and usage patterns to determine which cloud provider and reservation options best suit their cost optimization and resource management requirements.

4. How to Overcome the RI Resale Ban: Alternatives to Reserved Instances

1. AWS Savings Plans

When planning future purchases, one way to deal with the ban is to go for AWS Savings Plans instead.

A Savings Plan is a pricing scheme that offers a discount on On-Demand instances in exchange for committing to one or three years of use, where you set your daily spending limit. Until that point, all computing usage will be available at a reduced cost. When you exceed your limit,

AWS charges you the normal on-demand price.

EC2 Savings Plan vs. Compute Savings Plan

The EC2 Instance Savings Plan offers up to 72% price reductions on EC2 instances, and you can choose the plan's size, OS, and tenancy. AWS also provides Compute Savings Plans, which give a similar discount amount (66% versus 72%) and include choices such as family, region, operating system, tenancy, and even individual compute services. EC2 Savings Plans eliminate the need for manual reservation procurement for specific instance families. This streamlines the process and reduces the need for coordination across multiple teams, minimizing the risk of errors.

2. Look Beyond Commitment for Cost Savings

Savings Plans can help you save money on AWS, but you're in charge of infrastructure optimization.

This is why picking the right size and compute instances is essential. If you manage a large cloud environment, you'll need a system that automates cost optimization activities like rightsizing, autoscaling, and instance type selection.

It takes time to determine which resources are running, which families control them, and whose teams own them. Trying to make sense of all the 500+ EC2 instances AWS offers is no walk in the park. It can take you many days or weeks to assess your inventory and use it to determine which instances to keep and which ones to eliminate.

AWS policy change means commitment plans are now at a 100% consumption lock-in. Using them for workloads with a non-steady demand is risky. Use an automation solution that matches resources to real-time application demand with changing demand.

5. Capacity Reservations

Capacity Reservations enable you to reserve computing capacity for your Amazon EC2 instances in a specific availability zone for any duration. They mitigate the risk of being unable to get On-Demand capacity in case of capacity constraints. Suppose you have strict capacity requirements and are running business-critical workloads that require a certain level of long-term or short-term capacity assurance. In that case, we recommend creating a Capacity Reservation to ensure that you always have access to Amazon EC2 capacity when you need it for as long as you need it. You can create Capacity Reservations anytime without entering into a one-year or three-year term commitment. Once the capacity reservation is provisioned in your

account, the capacity becomes available, and billing starts. When you no longer need the capacity assurance, cancel the Capacity Reservation to release the capacity and stop incurring charges. You can also use the billing discounts offered by Savings Plans and Regional Reserved Instances to reduce the cost of a Capacity Reservation.

When you create a Capacity Reservation, you specify:

- The Availability Zone in which to reserve the capacity
- The number of instances for which to reserve capacity The instance attributes, including the
- The instance attributes, including the

Capacity Reservations can only be used by instances that match their attributes. By default, they are automatically used by running instances that match the attributes. If you don't have any running instances that match the attributes of the Capacity Reservation, they remain unused until you launch an instance with matching attributes.

Differences Among Capacity Reservations, Reserved Instances, and Savings Plans

The following table highlights key differences between Capacity Reservations, Reserved Instances, and Savings Plans:

| | Capacity Reservations | Zonal Reserved Instances | Regional Reserved Instances | Savings Plans |
|------------------|---|--|---|-----------------------|
| Term | No commitment required. Can be created and cancelled as needed. | Requires a fixed one-year or three-year commitment | | |
| Capacity benefit | Capacity reserved in a specific Availability Zone. | | | No capacity reserved. |
| Billing discount | No billing discount. | Provides a billing discount. | | |
| Instance Limits | Your On-Demand Instance limits per Region apply | The default is 20 per Availability Zone. You can request a limit increase. | The default is 20 per Region. You can request a limit increase. | No limit. |

You can combine Capacity Reservations with Savings Plans or Regional Reserved Instances to receive a discount.

7. Pull Request (PR) to Purchase and Track RIs and SPs

Create a Pull Request (PR) in a code repository like GitHub or Bitbucket to track Reserved Instance (RI) and Savings Plan (SP) commitments and ensure accountability. The PR process typically involves updating documentation, creating or modifying scripts/tools for tracking, and implementing automated alerts or notifications. Here's a basic outline of steps you might take to create such a PR:

A. Update Documentation:

- Update the documentation to include information about tracking RI and SP commitments, including their importance, how they contribute to cost optimization, and the process for tracking and managing them.
- Include details about the types of reservations available, their terms, and any requirements or best practices for tracking them effectively.

B. Scripts/Tools for Tracking:

- Develop or modify scripts/tools to track RI and SP commitments. These scripts/tools could query the respective cloud provider's API to retrieve information about active reservations.
- Ensure the scripts/tools capture relevant details such as reservation IDs, instance types, terms, start and end dates, payment options, and associated costs or savings.

C. Automated Alerts/Notifications:

- Implement automated alerts or notifications to notify relevant stakeholders about upcoming expiration dates of reservations, changes in reservation status, or any discrepancies in tracked commitments.
- Configure alerts to be sent via email, messaging platforms, or integrated into existing monitoring systems for visibility and accountability.

D. Testing:

- Test the scripts/tools thoroughly to capture and report RI and SP commitments accurately. Verify that alerts and notifications are triggered correctly based on predefined criteria.
- Conduct testing in a staging environment to validate the tracking system's functionality and performance before deploying it to production.

E. Review and Approval:

- Submit the PR for review, ensuring that it includes all necessary changes to documentation, scripts/tools, and any associated resources.

- Collaborate with relevant stakeholders, such as finance, operations, and cloud governance teams, to gather feedback and approvals and address concerns.

F. Deployment:

- Once approved, merge the PR into the appropriate branch or repository to deploy the changes.
- Monitor the tracking system post-deployment to ensure that it functions as expected and continues to provide accurate and timely information about RI and SP commitments.

G. Ongoing Maintenance:

- Review and update the tracking system regularly to accommodate

changes in reservation terms, pricing models, or cloud provider APIs.

- Maintain documentation to reflect any updates or improvements made to the tracking process.
- By following these steps and creating a PR encompassing documentation update, script tool development, testing, review, deployment, and ongoing maintenance, you can establish a robust system for tracking RI and SP commitments and ensuring accountability within your organization.

CONCLUSION

Effectively planned and managed reservations can help you achieve significant discounts for AWS workloads that run on a predictable schedule. It's crucial to analyze your current AWS usage to select the right reservation attributes from the start and to devise a longer-term strategy for monitoring and managing your Reserved Instances. Using tools such as the AWS Compute Optimizer, AWS Cost and Usage report, and the Reserved Instance Utilization and Coverage reports in AWS Cost Explorer, you can examine your overall usage and discover opportunities for greater cost efficiencies.

Cloud cost optimization is not a one-time activity but an ongoing process, and cutting cloud expenses is not just the responsibility of IT but the entire organization. It is necessary to regularly review and assess your cloud costs to ensure that you are using the most cost-effective solutions for your needs. By

implementing the strategies discussed in this article, you can effectively manage your cloud costs and ensure your organization gets the most out of its cloud investment.

GCP discounts include sustained-use discounts and committed-use discounts. The former are applied automatically, and the latter are based on a 1-year or 3-year commitment. Sustained-use discounts are automatically applied to Compute Engine instances as long as they are used for over 25% of the month. Committed-use discounts have two types: resource-based and spend-based. The former is exclusive to Compute Engine instances, whereas the latter can apply to Cloud SQL, Cloud VMware Engine, Cloud Run, and Google Kubernetes Engine.

Spot VMs can also be used for stateless applications as they also offer significant savings. It is ideal for fault-tolerant and stateless workloads due to the unpredictable nature of pre-emption.

REFERENCES

- [1] AWS Documentation. "Amazon EC2 Reserved Instances." [Online]. Available: <https://aws.amazon.com/ec2/pricing/reserved-instances/>, Accessed Feb, 2024.
- [2] Abud, Luis. (2022, September 28). "Easily save money on your compute spend with Savings Plans." AWS Cloud Financial Management Blog. [Online]. Available: <https://aws.amazon.com/blogs/aws-cloud-financial-management/easily-save-money-on-your-compute-spend-with-savings-plans/>
- [3] Google Cloud Documentation. "Resource-based committed use discounts." [Online]. Available: <https://cloud.google.com/compute/docs/instances/signing-up-committed-use-discounts>, Accessed Feb, 2024.
- [4] Microsoft Azure documentation. (2024, January 16). "How saving plan discount is applied." [Online]. Available: <https://learn.microsoft.com/en-us/azure/cost-management-billing/savings-plan/discount-application>.
- [5] Google Cloud Documentation. "Sustained use discounts for Compute Engine." [Online]. Available: <https://cloud.google.com/compute/docs/sustained-use-discounts>, Accessed Feb, 2024.