Microsoft Azure purchasing models guidance

Pay-as-you-go

Overview

The Azure pay-as-you-go model allows you to use Azure services without any upfront commitment or long-term contracts. This model is available to anyone who wants to start using Azure services immediately and pay only for what they use.

Summary

This guidance provides an overview of your options for acquiring Azure subscriptions and services. It covers the various purchasing and payment options, such as Pay-as-you-go, Azure Reserved Instances, Azure Savings Plan, and others. You will find related information about pricing, billing, cancellation policies, reporting, and more to help you choose the right option for your organization.

	Pay-as-you-go	Reserved Instance (RI)	Savings Plan (ASP)	Pre-Purchase Plan (P3)	Azure Spot
Overview	Pay in arrears only for what you use.	Best discount but limited to specific resource for whole term.	Lower discounts than RI but can bundle across multiple resources/services.	Lower discounts than RI. Same discount applies to all resources for service.	Heavily discounted with dynamic pricing and possibility of eviction.
Applicability	All metered Azure resources and services	All metered Azure resources and services	Available for Compute resources only	Limited to only a few services currently	Available mostly for Compute resources
Discount	None: "List" price	Up to 72%	Up to 65%	Up to 40%	40% - 90%
Commitment Term	None	1-year, 3-year, or 5-year (where available) terms. Cancellation allowed up to $50K/ rolling 12 months.	1-year or 3-year terms	1-year or 3-year terms	None
Commitment Scope	None	Units of the specific resource per hour	Dollar per hour across resources and services	Dollar amount over term	None
Billing	Billed monthly in arrears	Billed upfront or monthly for term	Billed upfront for term or prepaid monthly	Billed upfront for term	Billed monthly in arrears
Drawdown	N/A	Use it or lose it	Use it or lose it	Full term available for commitment drawdown	N/A

How you are charged

- Charges are based on actual usage metered per second and billed hourly, with no upfront costs or termination fees. Depending on the specific service, you are charged for each minute of service used (for example, a virtual machine running for 6 minutes 45 seconds is billed for 6 minutes) or your total monthly usage (for example, the total Azure Blob Storage you used for the month).
- No prepayment or long-term commitment is required.
- Service prices are listed on the Azure pricing page and vary by region and by product.
- You can start using services immediately and pay at the end of the monthly billing period (hourly or monthly), providing flexibility to scale as needed.

Billing schedule

You receive a monthly invoice that details your organization's usage and charges for the billing period. The invoice includes a breakdown of the costs for each service used, making it easy for you to understand the charges.

Cancellation policy

With pay-as-you-go, you can cancel at any time. You just need to pay for any outstanding charges.

Reporting

You can track usage and costs via the Azure portal, which offers a centralized view of resources and billing information. Azure Cost Management provides insights for optimizing spending.

Example scenario 1: Pay-as-you-go model for Virtual Machines

Scenario: Contoso, who is a growing tech startup, needs to purchase Virtual Machines (VMs) to support their expanding operations. They are not sure about their usage requirements and would like to explore different Azure services.

Solution: Contoso can choose the pay-as-you-go model for its flexibility and scalability. This model allows them to pay only for the resources they use, without any upfront commitment.

Contoso uses a compute resource—B2ls v2 2/4 in West—for 100 VMs for 15 hours a day, every day for one month (30 days). The pay-as-you-go rate for this resource is $0.06/VM/hour. Based on their usage per hour, their total monthly bill will be:

100 VMs × $0.06/VM/hour × 15 hours/day × 30 days/month = $2,700 which will be paid at the end of the month.

Example scenario 2: Pay-as-you-go model for storage

Scenario: Contoso needs to purchase storage to handle large volumes of data generated by their analytics processes. They are not sure about their storage requirements and would like to explore different Azure storage services.

Solution: Contoso can choose the pay-as-you-go model for its cost-effectiveness and ease of scaling. This model allows them to pay only for the storage they use, without any long-term commitment.

Contoso uses a Blob Storage (Prem in West) resource for 2,000 GBs for 24 hours a day, 10 days in a month. The pay-as-you-go rate for this resource is $0.0003/GB/hour. Based on their usage per hour, their total monthly bill will be:
2,000 GBs × $0.0003/GB/hour × 24 hours/day × 10 days/month = $144, which will be paid at the end of the month.

How you are charged

- You can purchase Reserved Instances through a one-time upfront payment or monthly payments.
- You are charged on an hourly basis for the committed quantity, regardless of usage. For example, if you commit to 10 virtual machines (VMs) in East US, you will be charged for 10 VMs per hour for the entire term commitment.
- Any additional usage beyond the hourly commitment is billed at pay-as-you-go rates. For example, if you have a Reserved Instance for 10 VMs in US Northeast and use 8 VMs in an hour, the Reserved Instance covers the 8 VMs. If usage exceeds to 15 VMs, the additional 5 VMs are billed at Pay-as-you-go rates. This model offers savings for stable workloads while maintaining flexibility for extra usage.
- Unused capacity within a given hour cannot be carried forward to the next hour.
- Reserved Instance benefits automatically apply to usage that matches Reserved Instance attributes such as SKU, region, etc.

Resource commitment

All increases in service usage, such as adding more running compute instances or increasing storage capacity, are subject to the availability of resources.

Any quotas described in the limits, quotas, and constraints section do not guarantee sufficient resources will be available to meet those quotas.

If an Azure free account is not incurring charges or is inactive for more than 90 days, Microsoft reserves the right to remove any compute resources including hosted service(s), virtual machines, and websites. We will provide notice prior to removing any compute resources. You can avoid becoming inactive by doing either of the following every 90 days:

- Sign in to the Azure portal.
- Access your resources via our Azure APIs.

Microsoft reserves the right to modify or discontinue any free services or Azure free account at any time with or without notice.

Azure Reserved Instances

Overview
Azure Reserved Instances (RIs) offer up to 72% savings compared to pay-as-you-go pricing when you commit to specific Azure services in a set region for 1-year or 3-year terms.

Availability

You can check the availability of Reserved Instances on the Azure Portal. All the VMs are supported by a Reserved Instance.

How to choose the right reservation

Azure provides a Reserved Instance (RI) purchase recommendation by analyzing historical pay-as-you-go usage patterns to help customers identify the most cost-effective Reserved Instance size based on past and projected usage.

Cancellation policy

Exchange policies:

- You can exchange reservations of the same type (for example, VM for VM, not VM for SQL Server).
- Exchanges are treated as a refund and new purchase.
- Only reservation owners can exchange.
- Exchanges can occur after the Enterprise Agreement expires.
- The new reservation term must meet or exceed the remaining commitment.
- Exchanges are penalty-free, with no annual limits.

Refund policies:

- Refunds are based on the lower of purchase or current price.
- Refunds are limited to $50,000 in 12 months per billing profile.
- CSP refund limit is per customer.

Reporting

You can track your usage and costs via the Azure portal, which offers a centralized view of resources and billing information.

Example scenario 1: Reserved Instance model for Virtual Machines

Scenario: Contoso initially chose the pay-as-you-go model for its flexibility. They were paying $2,700 per month for a compute resource – B2ls v2 2/4 in West for 100 VMs used for 15 hours per day, 30 days per month, with no discounts. However, as their usage became more predictable, they are now looking into committing to specific usage to get discounts on the regular pay-as-you-go price.

Solution: When Contoso commits to the same resource on Reserved Instance (RI) for hourly usage, they can receive significant discount over pay-as-you-go pricing for that specific resource. Now with RI, the same usage of 100 VMs can receive a 54% discount on the pay-as-you-go rate. They commit to using 100 VMs for 24 hours a day, 30 days a month. The pay-as-you-go rate for this resource is $0.06/VM/hour, but with the RI discount, the rate is reduced to $0.027/VM/hour. Based on their usage per hour, their total monthly bill will be:

100 VMs × $0.027/VM/hour × 24 hours/day × 30 days/month = $1,944, which is paid upfront at the beginning of the term or prepaid monthly.

Example scenario 2: Overages with existing Reserved Instance model

If Contoso needs additional 50 VMs on top of their commitment for any hour, these get charged at the regular pay-as-you-go rate of $0.06/VM/hour.

For example, if they use the additional 50 VMs for 15 hours a day for 10 days in a month, the additional cost will be:

50 VMs × $0.06/VM/hour × 15 hours/day × 10 days/month = $450. This is additional cost on top of their RI commitment.

Example scenario 3: Under usage with existing Reserved Instance model

In case of under usage, even if Contoso has unused resources, they still pay what they have committed for their per hour usage in their term. So, if instead of using 100 VMs per hour, their usage drops to 50 VMs per hour, they still pay $1,944 per month throughout their commitment term.

Azure Savings Plan

Overview

The Azure Savings Plan is a flexible pricing model that allows you to save money on your Azure compute services. Savings Plan offers up to 65% savings compared to pay-as-you-go pricing when you commit to spend a minimum amount per hour across eligible services. Currently, Azure only offers Savings plan for compute products.

How you are charged

- When you commit to spending a set amount per hour, Azure uses that commitment to apply a discount to all qualifying resources up to the committed amount.

- Each hour, your eligible compute usage is discounted up to your commitment amount, and any usage beyond that is billed at the regular pay-as-you-go rate.
- Your Savings Plan discounts are applied automatically to any eligible usage that falls within your committed hourly spend, allowing you to benefit from discounted rates without needing to specify services or regions.

Capacity limitations/constraints

Azure Reserved Instances (RIs) are subject to the limitations and constraints in the table below.

Azure RI Limitation/Constraint	Description
Limited availability by region	RIs are subject to capacity availability in specific Azure regions. If a region reaches capacity for a particular VM size or SKU, you may not be able to purchase RIs for that configuration.
VM sizes and families	Reservations are tied to specific VM sizes and families. If you want to reserve a different size or family, you'll need to purchase a new reservation.
Scope restrictions	RIs can be scoped to a specific subscription, resource group, or management group. This can limit flexibility if you need to utilize the reserved capacity across different scopes.
Service Level Agreement (SLA) and commitment	Microsoft does not always guarantee specific capacity availability, especially during peak usage times.
Instance types not covered	Not all Azure services or instance types of support RIs. Make sure to verify which services are eligible for reserved capacity.
Scaling limits	If your workload exceeds the reserved instances, you may need to rely on pay-as-you-go pricing for the additional capacity, which could lead to higher costs.
Regional flexibility	Some reservations allow for cross-region flexibility, but this may not apply to all RIs, limiting where you can use the reserved capacity.

Using Saving Plans with Reserved Instances

You can combine Savings Plans with Reserved Instances. This flexible application of discounts across multiple services and regions allows you to dynamically meet your usage needs while maintaining predictable, reduced costs within their commitment.

Reserved Instance discounts are applied before Savings Plans discounts. If you have an existing Reserved Instance commitment, your Reserved Instance discount is applied, and then any additional usage beyond your Reserved Instance commitment is discounted until you reach your hourly Savings Plan commitment. For example, if the commitment is $10 per hour, any eligible services used within that $10 cap receive the discounted rate. If usage in a particular hour exceeds the commitment, the excess usage is charged at regular pay-as-you-go rates.

Cancellation policy

You cannot cancel or exchange Saving Plans after you've purchased.

Reporting

You can track usage and costs via the Azure portal, which offers a centralized view of resources and billing information. Azure Cost Management provides insights for optimizing spending.

Example scenario 1: Savings Plan model for Virtual Machines

Scenario: Contoso initially chose the pay-as-you-go model and was paying $2,700 per month for 100 VMs used for 15 hours per day, 30 days per month. They then switched to the Reserved Instance (RI) model, paying $1,944 per month for the same usage, receiving a 54% discount on one specific resource. However, Contoso's requirements evolved, and they needed multiple VM instances. They are now looking for multiple resources based on their needs and are looking to package their needs for discounts.

Solution: With the Azure Savings Plan, Contoso can bundle multiple services for their hourly commitment, offering flexibility and cost savings across different resources. They commit to an hourly spend for their commitment term but receive a lesser discount compared to RI.

Contoso's requirement is for 100 VMs of one VM instance and 35 VMs of another instance. Their per-hour commitment comes to $5 per hour. Hence, their monthly spend on savings plan is going to be calculated as:

$5/hour × 24 hours/day × 30 days/month = $3,600

Now with Savings Plan commitment, Contoso can combine its resources for up to 49% discount. For the first VM instance priced at $0.06/VM/hour price gets discounted 46% to $0.031/VM/hour. Another VM instance priced at $0.11/VM/hour gets discounted 49% to $0.055/VM/hour.

With Savings plan, Contoso keeps getting the discounted rate until their hourly usage is within $5/hour. So, with 100 VMs of first instance and 35 VMs of second instance, Contoso's hourly usage is:

100 VMs × $0.031/hour = $3.1 per hour

35 VMs × $0.055/hour = $1.9 per hour, totaling to $5 per hour.

Example scenario 2: Overages with Savings Plan model for Virtual Machines

If Contoso's usage exceeds the hourly commitment of $5 their additional resources are charged on regular pay-as-you-rates. For example, if they use the additional 50 VMs for 15 hours a day for 10 days in a month, the additional cost will be:

50 VMs × $0.06/VM/hour × 15 hours/day × 10 days/month = $450. This is additional cost on top of their Savings Plan commitment.

Example scenario 3: Under usage with existing Savings Plan model

Even if Contoso's hourly spend on their resources is less than $5/hour (they have used less resources for any hour), they will be charged at $5/hour as per their Savings Plan commitment.

Pre-Purchase Plan

Overview
The Azure Pre-Purchase Plan offers flexibility in terms of resource types and customization. Customers can choose specific types of resources and customize their plan based on their needs. Customers pre-purchase their commitment for a 1 or 3-year term. This means they pay upfront for their commitment, which covers their estimated usage for the term. Currently Pre-Purchase Plan is only available for Azure Databricks, Microsoft Sentinel, and Microsoft Defender for Cloud.

- **Azure Databricks:** The Azure Databricks Unit (DBU) pre-purchase plan enables customers to pre-pay for Databricks compute for a one- or three-year term. This plan covers all Azure Databricks SKUs, including Premium and Standard SKUs for Data Engineering Light, Data Engineering, and Data Analytics.
- **Microsoft Sentinel:** The Sentinel Pre-Purchase Plan allows customers to pre-pay for Sentinel services, providing a uniform discount across all Sentinel meters. This plan incentivizes higher commitment spending with a tiered discount model.
- **Microsoft Defender for Cloud:** The Pre-Purchase Plan for Microsoft Defender for Cloud allows customers to pre-pay for Defender for Cloud services, providing cost savings through discounted tiers.

How you are charged

- You choose the specific amount you want to pay upfront, and that amount is credited toward your usage of Azure services during the ensuing term.
- As you use Azure services, any charges are drawn down from your prepaid credit based on the discounted rate.
- If your usage exceeds your prepaid amount, additional consumption is charged at standard pay-as-you-go rates.
- Any unused amount is forfeited at the end of the term.

Billing schedule

Under Azure pre-purchase plan, you will only be billed for consumption that exceeds your pre-paid credit amount. In this case, you will receive an invoice at the end of your term.

Cancellation policy

Once a commitment is made and payment is processed, pre-purchase plans are non-refundable.

Reporting

You can track usage and costs via the Azure portal, which offers a centralized view of resources and billing information. Azure Cost Management provides insights for optimizing spending.

Example scenario 1: Pre-Purchase Plan for Microsoft Defender for Cloud

Scenario: Contoso has several long-term projects that require sustained use of Azure resources. For instance, they have a data analytics project that involves continuous data processing and storage. They also have a machine learning project that requires consistent compute resources for training models. Given these predictable and consistent usage patterns, Contoso is looking for a way to optimize their cloud spending and manage their budget more effectively.

Solution: The Pre-Purchase Plan offers Contoso the flexibility to choose specific types of resources and customize their plan based on their needs. By pre-purchasing their commitment for a 1- or 3-year term, they can lock in lower rates and save on costs. This plan allows them to pay upfront for their estimated usage, providing cost stability and predictability.

For example, Contoso is using Microsoft Defender for Cloud servers. On the pay-as-you-go model, they are paying $0.02 per MDC unit per hour. For 100 MDCs used for 15 hours a day for 10 days, they would pay $300 at the end of the month. However, if they can predict their usage for a 1-year term, they can pre-purchase their resources at a discounted rate of $0.018 per MDC unit per hour. For the same usage, they would prepay $9,855 for the entire year, benefiting from the discounted rate.

Azure Spot

Overview
Azure Spot Instances offer you a cost-effective way to run interruptible workloads, with savings of 40-90% compared to pay-as-you-go pricing. Spot Instances leverage surplus Azure capacity, allowing you to access compute power at highly discounted rates, ideal for short-term, flexible workloads like batch processing, simulations, or testing.

How you are charged

- Spot Instance pricing is dynamic and may change at any time based on market demand and availability. Prices are published monthly but are not disclosed in advance, so you should monitor usage and set price thresholds. See the current Spot prices.
- You can set a maximum price, or "Pmax," indicating the highest amount you are willing to pay per hour. This enables cost control, as instances will only run if the market price stays within the set threshold.

Billing schedule

- Spot Instances are billed separately from regular Azure VMs, with their own metering and billing cycle.
- Spot billing follows the standard Azure monthly cycle, providing transparency on costs related to these dynamically priced instances.

Eviction policy

Customers can be evicted from a Spot Instance if the price exceeds their set Pmax or if Azure needs the capacity back. Eviction notifications are sent 30 seconds before termination, but customers must sign up for notifications to receive them; otherwise, evictions may occur immediately with as little as no notice time.

Example scenario 1: Azure Spot for short-term flexible workloads

Scenario: Contoso, a growing tech startup, is looking for a cost-effective solution for their short-term, flexible workloads such as batch processing, simulations, and testing. These workloads do not require continuous availability and can tolerate interruptions. Contoso initially used the pay-as-you-go model, paying $0.06 per VM per hour. For using 200 VMs for 2 hours in a day, they would pay $24. However, they are exploring ways to reduce costs for these specific workloads.

Solution: With Azure Spot, Contoso can take advantage of significant discounts by using unused Azure capacity. The Azure Spot offers about a 55% discount, dropping the price to $0.025 per VM per hour. For using 200 VMs for 2 hours in a day, the customer would now pay:

200 VMs × $0.025/VM/hour × 2 hours/day = $10

Therefore, Contoso will be billed $10 for the same usage, compared to $24 with the standard pay-as-you-go pricing.

Reserved capacity

You can reserve Azure compute capacity in a specific region or availability zone for any duration of time. This feature ensures that the reserved capacity is always available to your organization, even during periods of high demand or unexpected events.

Benefits of reserved capacity

- **Guaranteed availability:** After deployment, capacity is reserved for your use and is always available within the scope of applicable service-level agreements (SLAs). By reserving capacity, you can ensure Azure compute resources are consistently available regardless of fluctuating demand, helping minimize the risk of resource shortages for uninterrupted workloads.
- **Flexible reservation duration with no term commitment**: Reserved Capacity does not require a fixed 1-year or 3-year commitment. Capacity can be deployed and deleted at any time, and customers can reserve capacity for the exact duration needed, providing flexibility for projects or workloads with variable timelines.
- **Region- or zone-specific resources:** Customers can specify the region or availability zone where dedicated capacity is required, ensuring resources are strategically located to meet performance, operational, or regulatory requirements.
- **Cost optimization with Reserved Instances**: Reserved Capacity can be automatically combined with Reserved Instances to take advantage of term-commitment discounts where applicable.

For more details, refer to the On-demand capacity reservation in Azure article.

Ideal cases

Reserved Capacity is particularly useful for:

- **Mission-critical applications:** Ensures that high-priority applications always have the compute power they need.
- **Disaster recovery:** Supports continuity planning by guaranteeing capacity is available in specific zones during an outage.
- **Consistent workloads:** Suitable for steady, ongoing workloads that can't risk downtime due to capacity limitations.

Pricing and billing

Capacity reservations are priced at the same rate as the underlying VM size. For example, if you create a reservation for 10 D2s_v3 VMs, you start getting billed for 10 D2s_v3 VMs, even if the reservation isn't being used.

If you then deploy a D2s_v3 VM and specify the reservation property, the capacity reservation gets used. After the VM is in use, you pay for only the VM and not the capacity reservation. Let's say you deploy six D2s_v3 VMs against the previously mentioned capacity reservation. You see a bill for six D2s_v3 VMs and four unused capacity reservations, both charged at the same rate as a D2s_v3 VM.

Both used and unused capacity reservations are eligible for Savings Plan and Reserved Instances term commitment discounts. The remaining eight D2s_v3 are billed normally. The term commitment discounts could be applied on either the VM or the unused capacity reservation.

Microsoft Azure Consumption Commitment (MACC)

Overview
Microsoft Azure Consumption Commitment (MACC) is a flexible purchasing program for enterprise customers who commit to spending a minimum amount on Azure services over a specific term, usually one, two, or three years.

How it works

With MACC, an organization signs a contract with Microsoft agreeing to spend a certain amount on Azure services within the specified term. The committed spend amount applies across nearly all Azure services, giving customers the freedom to choose services based on their needs as they evolve. Unlike traditional reserved capacity models that lock customers into specific resources, MACC lets them apply their commitment to a broad range of services, making it suitable for companies with changing workloads or complex needs.

Billing and spending commitment

By committing to a minimum spend, customers can unlock special discounts and cost-saving benefits. MACC rates are often better than standard pay-as-you-go prices, helping organizations manage their budgets more predictably.

Azure keeps track of the customer's consumption monthly and deducts it from the agreed commitment. Each month, customers receive a bill showing their total Azure usage, remaining commitment balance, and any additional costs if they exceed their commitment level.

MACC provides clear reporting so customers can monitor their Azure spending, assess which services are drawing the most resources, and identify areas for optimization. This transparency helps organizations avoid unexpected costs and stay within their committed amount.

Benefits

- MACC provides cost savings through discounted rates on a variety of services.
- Customers can shift spending across Azure services as needs change without losing the benefits of their commitment.
- Organizations can plan their budgets over a multi-year term, making it easier to allocate resources and manage costs.

Azure Commitment Discount (ACD)

Overview
The Azure Commitment Discount (ACD) is a discount provided to customers who commit to a Microsoft Azure Consumption Commitment (MACC) on Azure services over a specified term.

How it works

- The ACD allows customers to receive discounts on their Azure services by committing to a specific amount of usage over a 1-year or 3-year term.

- The discount rates vary depending on the commitment term and the services included in the agreement.
- Customers can benefit from lower rates compared to pay-as-you-go pricing, and the discount is applied to the services used within the commitment period.
- The ACD is part of the broader Microsoft Azure Consumption Commitment (MACC) program, which includes various incentives and discounts for long-term commitments.

Limits, quotas, and constraints

Below is an overview of limits, quotas, and constraints. See the Azure subscription and service limits, quotas, and constraints article for additional details. If you need to exceed these limits, contact Azure customer service.

Limits

Limits refer to the maximum capacity or usage that a particular Azure service can handle. These limits can be either adjustable or fixed:

- **Adjustable limits:** Some services have limits that can be increased upon request. For example, the number of virtual machines (VMs) you can deploy in a region can be increased by submitting a support request
- **Fixed limits:** Other services have fixed limits that cannot be changed. For example, the maximum size of a single Azure SQL Database is fixed and cannot be increased beyond a certain point.

Quotas

Quotas are predefined thresholds that limit the amount of resources you can consume within a specific time frame. They help manage resource allocation and prevent unexpected costs:

- **Subscription quotas:** These quotas apply to the entire Azure subscription. For example, the number of virtual networks or storage accounts you can create within a subscription.
- **Service quotas:** These quotas apply to specific services within Azure. For example, the number of read and write operations per hour for the Azure Resource Manager API.

Constraints

Constraints are conditions or rules that govern how resources can be used or deployed. They ensure that resources are used efficiently and within the capacity of the Azure infrastructure:

- **Regional constraints:** Some limits and quotas are managed at a regional level. For example, the number of virtual CPUs (vCPUs) you can use in a specific region.
- **Resource constraints:** These constraints apply to individual resources. For example, the maximum number of resources you can deploy within a single resource group.

Reporting

Reporting for Azure Reserved Instances (RIs) involves tracking and analyzing the usage and cost benefits of purchasing RIs compared to pay-as-you-go pricing. Below is breakdown of how it works.

Tracking usage

Customers can view their Reserved Instance utilization in the Azure Portal under "Cost Management + Billing." These reports provide insights into how many hours of your Reserved Instance were used versus how many were available.

Utilization reports

- **Metrics:** Key metrics include utilization percentage (the ratio of used hours to reserved hours) and savings compared to pay-as-you-go rates.
- **Daily/monthly reports:** Customers can generate reports to review usage patterns and determine the effectiveness of their Reserved Instance investments.

Cost management

- **Cost analysis:** Use Azure Cost Management tools to analyze costs and visualize savings from RIs over time.

- **Budgeting:** Set budgets to monitor spending and compare it with expected savings from RIs.

Exporting data

- **CSV export:** Users can export their usage and cost data as CSV files for deeper analysis in tools like Excel.
- **APIs:** Azure provides APIs to programmatically access usage and billing data for custom reporting solutions.

Recommendations

Azure Advisor can provide recommendations on RIs based on historical usage patterns to optimize future reservations.

Managing reservations

You can adjust your reservations based on changing needs, including the ability to exchange or cancel RIs, subject to specific terms.

Compliance and governance

Ensure your usage aligns with organizational policies and compliance requirements, tracking how RIs impact overall cloud expenditure.