



AWS Whitepaper

How AWS Pricing Works



How AWS Pricing Works: AWS Whitepaper

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

.....	v
Abstract and introduction	1
Introduction	1
Are you Well-Architected?	1
Key principles	2
Understand the fundamentals of pricing	2
Start early with cost optimization	2
Maximize the power of flexibility	3
Use the right pricing model for the job	3
AWS Pricing/TCO Tools	4
AWS Pricing Calculator	4
Migration Evaluator	4
AWS Cost Optimization	6
Choose the right pricing models	6
Use RIs to reduce Amazon RDS, Amazon Redshift, Amazon ElastiCache, and Amazon	
OpenSearch Service costs	6
Amazon EC2 Cost Savings	6
Match capacity with demand	6
Identify Amazon EC2 instances with low-utilization, and reduce cost by stopping or	
rightsizing.	6
Identify Amazon RDS and Amazon Redshift instances with low utilization and reduce cost	
by stopping (RDS) and pausing (Redshift).	7
Analyze DynamoDB usage and reduce cost by leveraging AutoScaling or on-demand.	7
Implement processes to identify resource waste	7
Identify Amazon EBS volumes with low-utilization and reduce cost by snapshotting, then	
deleting them	7
Analyze Amazon S3 usage and reduce cost by leveraging lower cost storage tiers	8
Review networking and reduce costs by deleting idle load balancers	8
Cost calculation examples	9
AWS Cloud cost calculation example	9
Architecture	9
Daily usage profile	10
Amazon EC2 cost breakdown	11
Hybrid cloud cost calculation example	13

Hybrid architecture description	13
Hybrid cloud components selection	14
Hybrid cloud architecture cost breakdown	16
Conclusion	18
Contributors	19
Further reading	20
AWS Glossary	21
Document revisions	22
Notices	23

This whitepaper is for historical reference only. Some content might be outdated and some links might not be available.

How AWS Pricing Works

Publication date: **February 24, 2023** ([Document revisions](#))

Amazon Web Services (AWS) helps you move faster, reduce IT costs, and attain global scale through a broad set of global compute, storage, database, analytics, application, and deployment services. One of the main benefits of cloud services is the ability it gives you to optimize costs to match your needs, even as those needs change over time.

Introduction

AWS has the services to help you build sophisticated applications with increased flexibility, scalability, and reliability. Whether you're looking for compute power, database storage, content delivery, or other functionality, with AWS you pay only for the individual services you need, for as long as you use them, without complex licensing. AWS offers you a variety of pricing models for over 160 cloud services. You only pay for the services you consume, and once you stop using them, there are no additional costs or termination fees. This whitepaper provides an overview of how AWS pricing works across some of the most widely used services. The latest pricing information for each AWS service is available at: [AWS Pricing](#).

Are you Well-Architected?

The [AWS Well-Architected Framework](#) helps you understand the pros and cons of the decisions you make when building systems in the cloud. The six pillars of the Framework allow you to learn architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems. Using the [AWS Well-Architected Tool](#), available at no charge in the [AWS Management Console](#), you can review your workloads against these best practices by answering a set of questions for each pillar.

For more expert guidance and best practices for your cloud architecture—reference architecture deployments, diagrams, and whitepapers—refer to the [AWS Architecture Center](#).

Key principles

Although pricing models vary across services, it's worthwhile to review key principles and best practices that are broadly applicable.

Understand the fundamentals of pricing

There are three fundamental drivers of cost with AWS: compute, storage, and outbound data transfer. These characteristics vary somewhat, depending on the AWS product and pricing model you choose.

In most cases, there is no charge for inbound data transfer or for data transfer between other AWS services within the same Region. There are some exceptions, so be sure to verify data transfer rates before beginning. Outbound data transfer is aggregated across services and then charged at the outbound data transfer rate. This charge appears on the monthly statement as *AWS Data Transfer Out*. The more data you transfer, the less you pay per GB. For compute resources, you pay by the hour or by the second from the time you launch a resource until the time you stop or terminate it, unless you have made a reservation for which the cost is agreed upon beforehand. For data storage and transfer, you typically pay per GB.

Except as otherwise noted, AWS prices are exclusive of applicable taxes and duties, including value-added tax (VAT) and sales tax. For customers with a Japanese billing address, use of AWS is subject to Japanese Consumption Tax. For more information, see [Amazon Web Services Consumption Tax FAQ](#).

Start early with cost optimization

The cloud allows you to trade fixed expenses (such as data centers and physical servers) for variable expenses, and only pay for IT as you consume it. And because of the economies of scale, the variable expenses are much lower than what you would pay to do it yourself. Whether you started in the cloud or you are just starting your migration journey to the cloud, AWS has a set of solutions to help you manage and optimize your spend. This includes services, tools, and resources to organize and track cost and usage data, enhance control through consolidated billing and access permission, enable better planning through budgeting and forecasts, and further lower cost with resources and pricing optimizations. To learn how you can optimize and save costs today, visit [Optimize and Save your IT costs](#).

Maximize the power of flexibility

AWS services are priced independently and transparently, and available on-demand, so you can choose and pay for exactly what you need. You may also choose to save money through a reservation model. By paying for services on an as-needed basis, you can redirect your focus to innovation and invention, reducing procurement complexity and enabling your business to be fully elastic.

One of the key advantages of cloud-based resources is that you don't pay for them when they're not running. By turning off instances you don't use, you can reduce costs by 70 percent or more compared to using them 24/7. This enables you to be cost efficient and, at the same time, have all the power you need when workloads are active.

Use the right pricing model for the job

AWS offers several pricing models depending on product. These include:

- **On-Demand Instances** let you pay for compute or database capacity by the hour or second (minimum of 60 seconds) depending on which instances you run, with no long-term commitments or upfront payments.
- The **Savings Plans** flexible pricing model offers low prices on Amazon Elastic Compute Cloud (Amazon EC2), Amazon SageMaker AI, AWS Lambda, and AWS Fargate usage in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one or three-year term.
- The **Spot Instance** Amazon EC2 pricing mechanism lets you request spare computing capacity with no upfront commitment and at discounted hourly rate (up to 90 percent off the on-demand price).
- **Reservations** provide you with the ability to receive a greater discount (up to 75 percent) by paying for capacity ahead of time. For more details, see the [AWS Cost Optimization](#) section.

AWS Pricing/TCO Tools

To get the most out of your estimates, you should have a good idea of your basic requirements. For example, if you're going to try Amazon EC2, it might help if you know what kind of operating system you need, what your memory requirements are, and how much I/O you need. You should also decide whether you need storage, if you're going to run a database, and how long you intend to use the servers. You don't need to make these decisions before generating an estimate; you can play around with the service configuration and parameters to see which options fit your use case and budget best. For more information about AWS service pricing, see [AWS Pricing](#).

AWS offers free pricing and migration tools for you to use. If the workload details and services to be used are identified, AWS Pricing Calculator can help with calculating the total cost of ownership. Migration Evaluator helps with inventorying your existing environment, identifying workload information, and designing and planning your AWS migration.

AWS Pricing Calculator

AWS Pricing Calculator is a web-based service that you can use to create cost estimates to suit your AWS use cases. This service is useful both for people who have never used AWS and for those who want to reorganize or expand their usage.

AWS Pricing Calculator allows you to explore AWS services based on your use cases and create a cost estimate. You can model your solutions before building them, explore the price points and calculations behind your estimate, and find the available instance types and contract terms that meet your needs. This enables you to make informed decisions about using AWS. You can plan your AWS costs and usage or price out setting up a new set of instances and services.

AWS Pricing Calculator is free for use and provides an estimate of your AWS fees and charges (not including taxes). AWS Pricing Calculator provides pricing details for your information only. AWS Pricing Calculator provides a console interface at [AWS Pricing Calculator](#).

Migration Evaluator

Migration Evaluator (formerly TSO Logic) is a complimentary service to create data-driven business cases for AWS Cloud planning and migration.

Creating business cases on your own can be a time-consuming process and does not always identify the most cost-effective deployment and purchasing options. Migration Evaluator quickly

provides a business case to make sound AWS planning and migration decisions. With Migration Evaluator, your organization gets access to AWS expertise, visibility into multiple cost-effective cloud migration scenarios, and insights on reusing existing software licensing to further reduce costs.

A business case is the first step in the AWS migration journey. Beginning with on-premises inventory discovery, you can choose to upload exports from third-party tools or install a complimentary agentless collector to monitor Windows, Linux, and SQL Server footprints. As part of a white-glove experience, Migration Evaluator includes a team of program managers and solution architects to capture your migration objective and use analytics to narrow down the subset of migration patterns best suited to your business needs. The results are captured in a transparent business case which aligns business and technology stakeholders to provide a prescriptive next step in your migration journey.

Migration Evaluator service analyzes an enterprise's compute footprint, including server configuration, utilization, annual costs to operate, eligibility for bring-your-own-license, and hundreds of other parameters. It then statistically models utilization patterns, matching each workload with optimized placements in EC2 and Amazon Elastic Block Store (Amazon EBS). Finally, it outputs a business case with a comparison of the current-state against multiple future-state configurations showing the flexibility of AWS.

For more information, see [Migration Evaluator](#).

AWS Cost Optimization

AWS enables you to take control of cost and continuously optimize your spend, while building modern, scalable applications to meet your needs. AWS's breadth of services and pricing options offer the flexibility to effectively manage your costs and still keep the performance and capacity you require. AWS is dedicated to helping customers achieve the highest savings potential. Get started with the steps below that will have an immediate impact on your bill today.

Choose the right pricing models

Use RIs to reduce Amazon RDS, Amazon Redshift, Amazon ElastiCache, and Amazon OpenSearch Service costs

For certain services like Amazon EC2 and Amazon RDS, you can invest in reserved capacity.

With [Reserved Instances](#), you can save up to 72 percent over the equivalent on-demand capacity. RIs are available in three options: All up-front (AURI), partial up-front (PURI), and no upfront payments (NURI). Use the recommendations provided in AWS Cost Explorer RI purchase recommendations, which is based on your Amazon RDS, Amazon Redshift, ElastiCache, and OpenSearch Service usage.

Amazon EC2 Cost Savings

Use Amazon [Spot Instances](#) to reduce Amazon EC2 costs, use Compute [Savings Plans](#) to reduce Amazon EC2, Fargate, and Lambda costs, and use SageMaker [Savings Plans](#) to reduce SageMaker AI costs.

Match capacity with demand

Identify Amazon EC2 instances with low-utilization, and reduce cost by stopping or rightsizing.

Use [AWS Cost Explorer Resource Optimization](#) to get a report of Amazon EC2 instances that are either idle or have low utilization. You can reduce costs by either stopping or downsizing these instances. Use [AWS Instance Scheduler](#) to automatically stop instances. Use [AWS Operations](#)

[Conductor](#) to automatically resize the Amazon EC2 instances (based on the recommendations report from Cost Explorer).

Identify Amazon RDS and Amazon Redshift instances with low utilization and reduce cost by stopping (RDS) and pausing (Redshift).

Use the Trusted Advisor Amazon [RDS Idle DB instances check](#) to identify DB instances which have not had any connection over the last seven days. To reduce costs, stop these DB instances using the automation steps described here: [Implementing DB Instance Stop and Start in Amazon RDS](#). For Redshift, use the Trusted Advisor Underutilized [Redshift clusters check](#) to identify clusters which have had no connections for the last seven days, and less than 5 percent cluster wide average CPU utilization for 99 percent of the last seven days. To reduce costs, pause these clusters using the steps in: [Lower your costs with the new pause and resume actions on Amazon Redshift](#).

Analyze DynamoDB usage and reduce cost by leveraging AutoScaling or on-demand.

Analyze your DynamoDB usage by monitoring two metrics, ConsumedReadCapacityUnits and ConsumedWriteCapacityUnits, in CloudWatch. To automatically scale (in and out) your DynamoDB table, use the AutoScaling feature. Using the steps at [Enabling DynamoDB auto scaling on existing tables](#), you can enable AutoScaling on your existing tables. Alternately, you can also use the on-demand option. This option allows you to pay-per-request for read and write requests so that you only pay for what you use, making it easy to balance costs and performance.

Implement processes to identify resource waste

Identify Amazon EBS volumes with low-utilization and reduce cost by snapshotting, then deleting them

Amazon EBS volumes that have very low activity (less than one IOPS per day) over a period of seven days indicate that they are probably not in use. Identify these volumes using the Trusted Advisor Underutilized Amazon [EBS Volumes Check](#). To reduce costs, first snapshot the volume (in case you need it later), then delete these volumes. You can automate the creation of snapshots using the [Amazon Data Lifecycle Manager](#). Follow the steps at [Delete an Amazon EBS volume](#) to delete Amazon EBS volumes.

Analyze Amazon S3 usage and reduce cost by leveraging lower cost storage tiers

Use [Amazon S3 analytics](#) to analyze storage access patterns on the object data set for 30 days or longer. Amazon S3 Analytics makes recommendations for leveraging [S3 Infrequently Accessed](#) (S3 IA) to reduce costs. You can automate moving these objects into a lower cost storage tier using [lifecycle policies](#). Alternately, you can also use [S3 Intelligent-Tiering](#), which automatically analyzes and moves your objects to the appropriate storage tier.

Review networking and reduce costs by deleting idle load balancers

Use the Trusted Advisor Idle [Load Balancers check](#) to get a report of load balancers that have a RequestCount of less than 100 over the past seven days. Then use [Step 8: Delete your load balancer \(optional\)](#) to delete these load balancers to reduce costs. Additionally, use the steps provided in [Using AWS Cost Explorer to analyze data transfer costs](#) to review your data transfer costs using Cost Explorer.

Cost calculation examples

The following sections use the [AWS Pricing Calculator](#) to provide example cost calculations for two use cases.

Topics


- [AWS Cloud cost calculation example](#)
- [Hybrid cloud cost calculation example](#)

AWS Cloud cost calculation example

This example is a common use case of a dynamic website hosted on AWS using Amazon EC2, Amazon EC2 Auto Scaling, and Amazon RDS. The Amazon EC2 instance runs the web and application tiers, and Amazon EC2 Auto Scaling matches the number of instances to the traffic load. Amazon RDS uses one DB instance for its primary storage, and this DB instance is deployed across multiple Availability Zones.

Architecture

ELB balances traffic to the Amazon EC2 instances in an AWS Auto Scaling group, which adds or subtracts Amazon EC2 instances to match the load. Deploying Amazon RDS across multiple Availability Zones enhances data durability and availability. Amazon RDS provisions and maintains a standby in a different Availability Zone for automatic failover in the event of outages, planned or unplanned. The following illustration shows the example architecture for a dynamic website using Amazon EC2, Amazon EC2 Auto Scaling, and security groups to enforce least-privilege access to AWS infrastructure and selected architecture components, and one Amazon RDS database instance across multiple Availability Zones (Multi-AZ deployment). All these components are deployed into a single region and virtual private cloud (VPC). The VPC spans two availability zones to support failover scenarios. Route 53 Resolver is used to manage and route requests for one hosted zone towards the Elastic Load Balancer.



AWS Cloud deployment architecture

Daily usage profile

You can monitor daily usage for your application so that you can better estimate your costs. For instance, you can look at the daily pattern to figure out how your application handles traffic. For each hour, track how many hits you get on your website and how many instances are running, and then add up the total number of hits for that day.

$$\text{Hourly instance pattern} = (\text{hits per hour on website}) / (\text{number of instances})$$

Examine the number of Amazon EC2 instances that run each hour, and then take the average. You can use the number of hits per day and the average number of instances for your calculations.

$$\text{Daily profile} = \text{SUM}(\text{Hourly instance pattern}) / 24$$

Amazon EC2 cost breakdown

The following table shows the characteristics for Amazon EC2 used for this dynamic site in the US East (Northern Virginia) Region.

Characteristic	Estimated Usage	Description
Utilization	100%	All infrastructure components run 24 hours per day, seven days per week
Instance	t3a.xlarge	16 GB memory, 4 vCPU
Storage	Amazon EBS SSD gp2	One Amazon EBS volume per instance with 30 GB of storage per volume
Data backup	Daily Amazon EBS snapshots	One Amazon EBS volume per instance with 30 GB of storage per volume
Data transfer	Data in: 1 TB/month Data out: 1 TB/month	10% incremental change per day
Instance scale	4	On average per day, there are four instances running
Load Balancing	20 GB/hour	ELB is used 24 hours per day, seven days per week. It processes a total of 20 GB/hour (data in and data out)
Database	MySQL, db.m5.large instance with 8 GB memory, 2 vCPUs, 100 GB storage	Multi-AZ deployment with synchronous standby replica in a separate Availability Zone

The total cost for one month is the sum of the cost of the running services and data transfer out, minus the AWS Free Tier discount. We calculated the total cost using the [AWS Pricing Calculator](#).

Table: Cost breakdown

Service	Monthly	Annually	Configuration
ELB	\$87.60	\$1,051.20	Number of Network Load Balancers (one), Processed bytes per Network Load Balancer (NLB) for TCP (20 GB per hour)
Amazon EC2	\$439.16	\$5,269.92	Operating system (Linux), quantity (four), storage for each Amazon EC2 instance (General Purpose SSD (gp2)), storage amount (30 GB), instance type (t3a.xlarge)
Amazon Elastic IP address	\$0	\$0	Number of Amazon EC2 instances (one), Number of EIPs per instance (one)
Amazon RDS for MySQL	\$272.66	\$3,271.92	Quantity (one) db.m5.large, Storage for each Amazon RDS instance (General Purpose SSD (gp2)), storage amount (100 GB)
Amazon Route 53	\$183.00	\$2,196.00	Hosted Zones (1), Number of Elastic Network Interface


			s (2), Basic Checks Within AWS (0)
Amazon Virtual Private Cloud (Amazon VPC)	\$92.07	\$1,104.84	Data transfer cost, inbound (from: internet) 1 TB per month, outbound (to: internet) 1 TB per month, intra-Region 0 TB per month

Hybrid cloud cost calculation example

This example is a hybrid cloud use case of AWS Outposts rack, deployed on-premises and connected to AWS Cloud using AWS Direct Connect. Outposts extends the existing VPC from the selected AWS Region to the customer data center. Selected AWS services required to run on-premises (for example, Amazon Elastic Kubernetes Service (Amazon EKS)) are available on Outposts inside the Outpost Availability Zone, deployed inside a separate subnet.

Hybrid architecture description

The following example shows an Outpost deployment with distributed Amazon EKS service extending to on-premises environments.



AWS Outpost with Amazon EKS Control Plane and Data Plane Architecture

Architecture

- The Control Plane for Amazon EKS remains in the Region, which means in the case of Amazon EKS, the Kubernetes Primary node will stay in the Availability Zone deployed to the Region (not on the Outposts).
- The Amazon EKS worker nodes are deployed on the Outpost, controlled by a Primary node deployed in the Availability Zone.

Traffic Flow

- The EKS Control Plane Traffic between EKS, AWS metrics, and CloudWatch transits third-party networks (AWS Direct Connect/AWS Site-to-Site VPN to the AWS Region).
- The Application / Data Traffic is isolated from Control plane and distributed between Outposts and local network.
- Distribution of Amazon Machine Images (AMIs) (deployed on Outpost) is driven by central Amazon ECR in Region; however, all images are cached locally on the Outpost.

Load Balancers

- Application Load Balancer is supported on Outpost as the only local ELB available.
- The Network Load Balancer and Classic Load Balancer stay in the Region, but targets deployed at Outposts are supported (including Application Load Balancer).
- On-premises (inside corporate DC) Load Balancers (for example, F5 BIG IP, NetScaler) can be deployed and routed via Local Gateway (inside AWS Outpost).

Hybrid cloud components selection

Customers can choose from a range of pre-validated Outposts configurations (*Figure: Example Outposts architecture*) offering a mix of Amazon EC2 and Amazon EBS capacity designed to meet a variety of application needs. AWS can also work with customers to create a customized configuration designed for their unique application needs.

To identify the correct configuration, make sure to verify the deployment and operational parameters of the selected physical location for the AWS Outpost rack installation. The following

example represents a set of parameters highlighting facility, networking, and power requirements needed for location validation (selected parameter: example value):

Purchase Option: All Upfront

Term: 3 Years

Max on premises power capacity: 20kVA

Max weight: 2,500lb

Networking uplink speed: 100Gbps

Number of Racks: 1

Average Power Draw per Rack: 9.34

Constraint (power draw/weight): Power Draw

Total Outpost vCPU: 480

Total Outpost Memory: 2,496 GiB

In addition to minimum parameters, you should make deployment assumptions prior to any order to minimize the performance and security impact on existing infrastructure (selected question: example assumption).

Question: What is the speed of the uplink ports from your Outposts Network Devices (OND)?

Example answer: 40 or 100Gbps.

Question: How many uplinks per OND will you use to connect the AWS Outpost to your network?


Example answer: Four uplinks.

Question: How will the Outpost service link (the Outpost control plane) access AWS services?

Example answer: Service link will access AWS over a Direct Connect public VIF.

Question: Is there a firewall between Outposts and the Internet. Example answer: Yes

Considering these assumptions together with selected components will result in an architecture with a higher granularity of detail and will influence the overall cost of a hybrid cloud deployment (*Figure AWS Outpost with Amazon EKS Control Plane and Data Plane Architecture*).



Hybrid cloud architecture deployment example

Hybrid cloud architecture cost breakdown

Hybrid cloud costs include multiple layers and components deployed across the AWS cloud and on-premises location. When you use AWS Managed Services on Outposts, you are charged for the services based only on usage by instance-hour and not for the underlying Amazon EC2 instance and Amazon EBS storage.

Breakdown of these services is showcased in next sections for a three-year term with partial upfront, all upfront, and no upfront options (Amazon EC2 and Amazon EBS capacity). Price includes delivery, installation, servicing, and removal at the end of term—there is no additional charge.

Outpost rack charges (customized example)

Amazon EC2 Charges

- c5.24xlarge, 11 TB
- \$7,148.67 monthly;

- \$123,650.18 up front, \$3,434.73 monthly
- \$239,761.41 up front
- 1 m5.24xlarge, 11 TB
- \$7,359.69 monthly
- \$127,167.06 up front, \$3,532.42 monthly
- \$246,373.14 up front
- Amazon EBS
- 11 TB EBS tier is priced at \$0.30/GB monthly

Conclusion

Although the number and types of services offered by AWS have increased dramatically, our philosophy on pricing has not changed. You pay as you go, pay for what you use, pay less as you use more, and pay even less when you reserve capacity. All of these options empower AWS customers to choose their preferred pricing model and increase the flexibility of their cost strategy.

Projecting costs for a use case—for example, web application hosting—can be challenging because a solution typically uses multiple features across multiple AWS products. This means there are more factors and purchase options to consider.

The best way to estimate costs is to examine the fundamental characteristics for each AWS product, estimate your usage for each characteristic, and then map that usage to the prices posted on the website.

You can use the [AWS Pricing Calculator](#) to estimate your monthly bill. The calculator provides a per- service cost breakdown, as well as an aggregate monthly estimate. You can also use the calculator to see an estimation and breakdown of costs for common solutions.

Remember, you can get started with most AWS services at no cost using the [AWS Free Tier](#).

Contributors

Contributors to this document include:

- Vladimir Baranek, Principal Partner Solution Architect, Amazon Web Services
- Senthil Arumugam, Senior Partner Solutions Architect, Amazon Web Services
- Mihir Desai, Senior Partner Solutions Architect, Amazon Web Services
- Nick Worthington, Solutions Architect, Amazon Web Services
- Marc Weiss, Solutions Architect, Amazon Web Services

Further reading

For additional information, see:

- [AWS Pricing](#)
- [AWS Pricing Calculator](#)
- [AWS Free Tier](#)
- [AWS Cloud Financial Management](#)
- [AWS Cost and Usage Reports](#)
- [AWS Cloud Economics Center](#)

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.

Document revisions

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Whitepaper updated	Removed service-specific pricing, support plan information section, and Free Tier section.	December 18, 2024
Whitepaper updated	Removed Amazon Glacier Select pricing.	August 29, 2023
Whitepaper updated	Updated and added service pricing details, options, calculation, and examples.	February 24, 2023
Minor update	Fix non-inclusive language.	April 6, 2022
Whitepaper updated	Updated and added service pricing details, options, calculation, and examples.	October 30, 2020
Initial publication	Whitepaper first published.	June 1, 2018

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2023 Amazon Web Services, Inc. or its affiliates. All rights reserved.