

How AWS Pricing Works

June 2018



(Please consult <http://aws.amazon.com/whitepapers/> for the latest version of this paper)

Table of Contents

Introduction	4
Key Principles.....	4
Understand the fundamentals of pricing	4
Start early with cost optimization.....	4
Maximize the power of flexibility.....	5
Use the right pricing model for the job	5
Get started with the AWS Free Tier	6
Pricing details for individual services.....	6
Amazon Elastic Compute Cloud (Amazon EC2).....	6
AWS Lambda.....	9
Amazon Elastic Block Store (Amazon EBS).....	10
Amazon Simple Storage Service (Amazon S3).....	11
Amazon S3 Glacier	11
AWS Snowball	12
Amazon RDS	14
Amazon DynamoDB.....	15
Amazon CloudFront.....	16
Optimizing costs with reservations.....	17
Amazon EC2 Reserved Instances.....	17
Amazon DynamoDB Reserved Capacity	18
Amazon ElastiCache Reserved Nodes.....	18
Amazon RDS RIs	19
Amazon Redshift Reserved Nodes	19
Conclusion	19
Resources.....	20
Cost calculation example	21

©2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. 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.

Introduction

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.

AWS offers on-demand, pay-as-you-go, and reservation-based payment models, enabling you to obtain the best return on your investment for each specific use case. AWS services do not have complex dependencies or licensing requirements, so you can get exactly what you need to build innovative, cost-effective solutions using the latest technology. In this whitepaper, we'll provide an overview of how AWS pricing works across some of our most widely used services. The latest pricing information for each AWS service is available at <http://aws.amazon.com/pricing/>.

Key Principles

While 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 hourly from the time you launch a resource until the time you 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, our prices are exclusive of applicable taxes and duties, including VAT and sales tax. For customers with a Japanese billing address, use of AWS is subject to Japanese Consumption Tax. [Learn more](#).

Start early with cost optimization

Adopting cloud services is not just a technical evolution. It also requires changes to how organizations operate. As you move from IT being treated as a capital investment that happens periodically to a world where pricing is closely tied to efficient use of resources, it pays to understand what drives cloud pricing so you can build a strategy for optimizing it.

When it comes to understanding pricing and optimizing your costs, it's never too early to start. It's easiest to put cost visibility and control mechanisms in place before the environment grows large and complex. Managing cost-effectively from the start ensures that managing cloud investments doesn't become an obstruction as you grow and scale.

Learn more at the [AWS Cost Optimization page](#).

Maximize the power of flexibility

AWS services are priced independently and transparently, so you can choose and pay for exactly what you need and no more. No minimum commitments or long-term contracts are required unless you 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** means you pay for compute or database capacity with no long-term commitments or upfront payments.
- **Dedicated Instances** (available with [Amazon Elastic Compute Cloud](#) (Amazon EC2)) run in a virtual private cloud (VPC) on hardware that's dedicated to a single customer.
- **Spot Instances** are an Amazon EC2 pricing mechanism that lets you purchase spare computing capacity with no upfront commitment at discounted hourly rates.
- **Reservations** provide you with the ability to receive a greater discount, up to 75 percent, by paying for capacity ahead of time. More detail is provided in the section, "Optimizing costs with reservations." Reservation options include:
 - [Amazon EC2 Reserved Instances](#)
 - [Amazon DynamoDB Reserved Capacity](#)
 - [Amazon ElastiCache Reserved Nodes](#)
 - [Amazon Relational Database Service Reserved Instances](#)
 - [Amazon Redshift Reserved Nodes](#)

Get started with the AWS Free Tier

The AWS Free Tier enables you to gain free, hands-on experience with the AWS platform, products, and services. AWS Free Tier includes offers that expire 12 months after sign-up and others that never expire. This section lists some of the most commonly used AWS Free Tier offers. Many other services are available through AWS Free Tier, and terms and conditions apply. [Get full details here.](#)

The following free-tier offers are only available to new AWS customers, and are available for 12 months following your AWS sign-up date:

- [Amazon Elastic Compute Cloud \(Amazon EC2\)](#): 750 hours per month of Linux, RHEL, or SLES t2.micro instance usage or 750 hours per month of Windows t2.micro instance usage
- [Amazon Simple Storage Service \(Amazon S3\)](#): 5 GB of Amazon S3 standard storage, 20,000 Get Requests, and 2,000 Put Requests
- [Amazon Relational Database Service \(Amazon RDS\)](#): 750 hours of Amazon RDS Single-AZ db.t2.micro Instances for running MySQL, PostgreSQL, MariaDB, Oracle BYOL, or SQL Server (running SQL Server Express Edition); 20 GB of database storage; 10 million I/Os; and 20 GB of backup storage
- [Amazon CloudFront](#): 50 GB Data Transfer Out and 2,000,000 HTTP and HTTPS Requests each month

The following free-tier offers do not automatically expire at the end of your 12-month AWS Free Tier term and are available to all AWS customers:

- [AWS DynamoDB](#): Up to 200 million requests per month (25 [Write Capacity units](#) and 25 [Read Capacity units](#)); 25 GB of indexed data storage; 2.5 million read requests per month from [DynamoDB Streams](#); ability to deploy DynamoDB [Global Tables](#) in up to two AWS regions
- [Amazon Glacier](#): Retrieve up to 10 GB of your Amazon Glacier data per month for free
- [AWS Lambda](#): 1 million free requests per month; up to 3.2 million seconds of compute time per

Pricing details for individual services

Different types of services lend themselves to different pricing models. For example, Amazon EC2 pricing varies by instance type, while the Amazon Aurora database service includes charges for data input/output (I/O) and storage. This section provides an overview of pricing concepts and examples for a range of commonly used services. You can always find current price information for each AWS service at <http://aws.amazon.com/pricing/>.

Amazon Elastic Compute Cloud (Amazon EC2)

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers. The

simple web service interface of Amazon EC2 allows you to obtain and configure capacity with minimal friction with complete control of your computing resources.

Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change.

Pricing models for Amazon EC2

There are four ways to pay for Amazon EC2 instances: [On-Demand Instances](#), [Reserved Instances](#), [Spot Instances](#), and [Dedicated Hosts](#).

On-Demand Instances

With On-Demand Instances, you pay for compute capacity per hour or per second, depending on which instances you run. No longer-term commitments or upfront payments are needed. You can increase or decrease your compute capacity to meet the demands of your application and only pay the specified hourly rates for the instance you use. On-Demand Instances are recommended for:

- Users who prefer the low cost and flexibility of Amazon EC2 without upfront payment or long-term commitments
- Applications with short-term, spiky, or unpredictable workloads that cannot be interrupted
- Applications being developed or tested on Amazon EC2 for the first time

Spot Instances

[Amazon EC2 Spot Instances](#) allow you to request spare Amazon EC2 computing capacity for up to 90 percent off the On-Demand price. Spot Instances are recommended for:

- Applications that have flexible start and end times
- Applications that are only feasible at very low compute prices
- Users with urgent computing needs for a lot of additional capacity

The Spot price is set by Amazon EC2 and fluctuates periodically depending on the supply of, and demand for, Spot Instance capacity. If your maximum bid exceeds the current Spot price, your bid request is fulfilled, and your instances will run until either you choose to terminate them or the Spot price exceeds your maximum bid, whichever comes first. [Learn more about Spot Instances](#).

Reserved Instances

[Amazon EC2 Reserved Instances](#) provide you with a significant discount (up to 75 percent) compared to On-Demand Instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them.

For applications that have predictable usage, Reserved Instances can provide significant savings compared to On-Demand Instances. Reserved Instances are recommended for:

- Applications with steady-state usage
- Applications that require reserve capacity
- Customers who can commit to using EC2 over a 1- or 3-year term to reduce their total computing costs

Dedicated Hosts

A [Dedicated Host](#) is a physical EC2 server dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, including Windows Server, SQL Server, and SUSE Linux Enterprise Server (subject to your license terms). They can also help you meet compliance requirements.

Per-second billing

When Amazon EC2 was launched, the ability to use an instance for an hour, and to pay only for that hour, was big news. Today, many customers use Amazon EC2 to do a lot of work in a short time, sometimes minutes or even seconds. In 2017, we announced per-second billing for usage of Linux instances across On-Demand, Reserved, and Spot Instances. Similarly, provisioned storage for EBS volumes is billed in one-second increments.

Per-second billing saves money. It is particularly effective for resources that have periods of low and high usage such as development and testing, data processing, analytics, batch processing, and gaming applications. [Learn more about per-second billing.](#)

Estimating Amazon EC2 costs

When you begin to estimate the cost of using Amazon EC2, consider the following:

- **Clock hours of server time:** Resources incur charges when they are running—for example, from the time Amazon EC2 instances are launched until they are terminated, or from the time Elastic IPs are allocated until the time they are de-allocated.
- **Instance type:** Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for your applications. Each instance type includes at least one instance size, allowing you to scale your resources to the requirements of your target workload.
- **Pricing model:** With On-Demand Instances, you pay for compute capacity by the hour with no required minimum commitments. Reserved Instances give you the option to make a low one-time payment—or no payment at all—for each instance you want to reserve and in turn receive a significant discount on the hourly usage charge for that instance. With Spot Instances, you can bid for unused Amazon EC2 capacity. For additional information on how to save more money using Reserved Instances and Spot Instances, see the [AWS Cost Optimization page](#).
- **Number of instances:** You can provision multiple instances of your Amazon EC2 and Amazon EBS resources to handle peak loads.
- **Load balancing:** An Elastic Load Balancer can be used to distribute traffic among Amazon EC2 Instances. The number of hours the Elastic Load Balancer runs and the amount of data it processes contribute to the monthly cost.

- **Detailed monitoring:** You can use [Amazon CloudWatch](#) to monitor your EC2 instances. By default, basic monitoring is enabled. For a fixed monthly rate, you can opt for detailed monitoring, which includes seven preselected metrics recorded once a minute. Partial months are charged on an hourly pro rata basis, at a per instance-hour rate.
- **Auto Scaling:** Auto Scaling automatically adjusts the number of Amazon EC2 instances in your deployment according to conditions you define. This service is available at no additional charge beyond Amazon CloudWatch fees.
- **Elastic IP addresses:** You can have one Elastic IP (EIP) address associated with a running instance at no charge.
- **Operating systems and software packages:** Operating system prices are included in instance prices, unless you choose to bring your own licenses. There are no additional licensing costs to run the following commercial operating systems: [Red Hat Enterprise Linux](#), [SUSE Enterprise Linux](#), [Windows Server](#), and [Oracle Enterprise Linux](#). Additionally, AWS has made it easy for you by partnering with Microsoft, IBM, and several other vendors so you can run commercial software packages, such as [Microsoft SQL Server](#), on your Amazon EC2 Instances,. For commercial software packages AWS doesn't provide—such as nonstandard operating systems, Oracle Applications, Windows Server applications such as Microsoft SharePoint and Microsoft Exchange—you need to obtain a license from the vendors. You can also bring your existing license to the cloud through specific vendor programs such as [Microsoft License Mobility Through Software Assurance Program](#).

Get more information on [Amazon EC2 pricing](#).

AWS Lambda

AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume—there is no charge when your code is not running. With Lambda, you can run code for virtually any type of application or backend service—all with zero administration. Just upload your code and Lambda takes care of everything required to run and scale your code with high availability.

AWS Lambda pricing

With AWS Lambda, you pay only for what you use. You are charged based on the number of requests for your functions and the time it takes for your code to execute. Lambda registers a request each time it starts executing in response to an event notification or invoke call, including test invokes from the console. You are charged for the total number of requests across all your functions.

Duration is calculated from the time your code begins executing until it returns or otherwise terminates, rounded up to the nearest 100 milliseconds. The price depends on the amount of memory you allocate to your function.

Request pricing

- Free Tier: 1 million requests per month, 400,000 GB-seconds of compute time per month
- \$.20 per 1 million requests thereafter, or \$0.0000002 per request

Duration pricing

- 400,000 GB-seconds per month free, up to 3.2 million seconds of compute time
- \$0.00001667 for every GB-second used thereafter

Additional charges

You may incur additional charges if your Lambda function uses other AWS services or transfers data. For example, if your Lambda function reads and writes data to or from Amazon S3, you will be billed for the read/write requests and the data stored in Amazon S3.

Amazon Elastic Block Store (Amazon EBS)

[Amazon Elastic Block Store](#) (EBS) provides block level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance. They are analogous to virtual disks in the cloud. Amazon EBS provides two volume types:

- **SSD-backed volumes** are optimized for transactional workloads involving frequent read/write operations with small I/O size, where the dominant performance attribute is IOPS.
- **HDD-backed volumes** are optimized for large streaming workloads where throughput (measured in megabits per second) is a better performance measure than IOPS.

How Amazon EBS is priced

Amazon EBS pricing includes three factors:

- **Volumes:** Volume storage for all EBS volume types is charged by the amount of GB you provision per month, until you release the storage.
- **Snapshots:** Snapshot storage is based on the amount of space your data consumes in Amazon S3. Because Amazon EBS does not save empty blocks, it is likely that the snapshot size will be considerably less than your volume size. Copying EBS snapshots is charged based on the volume of data transferred across regions. For the first snapshot of a volume, Amazon EBS saves a full copy of your data to Amazon S3. For each incremental snapshot, only the changed part of your Amazon EBS volume is saved. After the snapshot is copied, standard EBS snapshot charges apply for storage in the destination region.
- **Data transfer:** Consider the amount of data transferred out of your application. Inbound data transfer is free, and outbound data transfer charges are tiered.

For more information, see the [Amazon EBS pricing page](#).

Amazon Simple Storage Service (Amazon S3)

Amazon S3 is object storage built to store and retrieve any amount of data from anywhere: websites, mobile apps, corporate applications, and data from IoT sensors or devices. It is designed to deliver 99.999999999 percent durability, and stores data for millions of applications used by market leaders in every industry. As with other AWS services, Amazon S3 provides the simplicity and cost-effectiveness of pay-as-you-go pricing.

Estimating Amazon S3 storage costs

With Amazon S3, you pay only for the storage you use, with no minimum fee. Prices are based on the location of your Amazon S3 bucket. When you begin to estimate the cost of Amazon S3, consider the following:

- **Storage class:** Standard Storage is designed to provide 99.999999999 percent durability and 99.99 percent availability. Standard – Infrequent Access (SIA) is a storage option within Amazon S3 you can use to reduce your costs by storing less-frequently accessed data at slightly lower levels of redundancy than the standard Amazon S3 storage. You can also use Amazon S3 Glacier storage for archiving data at very low costs. More information about Amazon S3 Glacier is provided below.
- **Storage:** Costs vary with number and size of objects stored in your Amazon S3 buckets as well as type of storage.
- **Requests:** The number and type of requests. GET requests incur charges at different rates than other requests, such as PUT and COPY requests.
- **Data transfer:** The amount of data transferred out of the Amazon S3 region.

Get more information about [Amazon S3 pricing](#). You can estimate your monthly bill using the [AWS Simple Monthly Calculator](#).

Amazon S3 Glacier

Amazon Glacier is a secure, durable, and extremely low-cost cloud storage service for data archiving and long-term backup. It is designed to deliver 99.999999999 percent durability, with comprehensive security and compliance capabilities that can help meet even the most stringent regulatory requirements. Amazon Glacier provides query-in-place functionality, allowing you to run powerful analytics directly on your archived data at rest.

Amazon Glacier provides low-cost, long-term storage

Starting at \$0.004 per GB per month, Amazon Glacier allows you to archive large amounts of data at a very low cost. You pay only for what you need, with no minimum commitments or upfront fees. Other factors determining pricing include requests and data transfers out of Amazon Glacier (incoming transfers are free).

Data access options

To keep costs low yet suitable for varying retrieval needs, Amazon Glacier provides [three options for access to archives](#) that span a few minutes to several hours.

	Expedited	Standard	Bulk
Data access time	1–5 minutes	3–5 hours	5–12 hours
Data retrievals	\$0.03 per GB	\$0.01 per GB	\$0.0025 per GB
Retrieval Requests	On-Demand: \$0.01 per request Provisioned: \$100 per Provisioned Capacity Unit	\$0.050 per 1,000 requests	\$0.025 per 1,000 requests

Storage and bandwidth include all file overhead

Rate tiers take into account your aggregate usage for Data Transfer Out to the internet across Amazon EC2, Amazon S3, Amazon Glacier, Amazon RDS, Amazon SimpleDB, Amazon SQS, Amazon SNS, Amazon DynamoDB, and AWS Storage Gateway.

Amazon Glacier Select pricing

Amazon Glacier Select allows queries to run directly on data stored in Amazon Glacier without having to retrieve the entire archive. Pricing for this feature is based on the total amount of data scanned, the amount of data returned by Amazon Glacier Select, and the number of Amazon Glacier Select requests initiated.

For more information, see the [Amazon Glacier pricing page](#).

AWS Snowball

AWS Snowball is a PB-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of the AWS cloud. Using Snowball addresses common challenges with large-scale data transfers, including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple, fast, and secure, and can be as little as one-fifth the cost of high-speed internet.

With AWS Snowball, you pay a service fee per data transfer job and the cost of shipping the appliance. Each job includes the use of a Snowball appliance for 10 days of onsite usage for free; there is a small charge for extra onsite days. Shipping days, including the day the appliance is received and the day it is shipped back to AWS, are not counted toward the 10 free days. Data transferred into AWS does not incur any data transfer fees, and [Standard Amazon S3 pricing](#) applies.

AWS Snowball pricing details

Service fee per job

Snowball 50 TB: \$200

Snowball 80 TB: \$250

Extra-day charge

The first 10 days of onsite usage are free. Each extra onsite day is \$15.

Data transfer

Data transfer in to Amazon S3 is free. Data transfer out of Amazon S3 is priced by region. For more information on AWS Snowball pricing, see the [AWS Snowball pricing page](#).

Amazon RDS

[Amazon RDS](#) is a web service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, so you can focus on your applications and business.

Estimating Amazon RDS costs

The factors that drive the costs of Amazon RDS include:

- **Clock hours of server time:** Resources incur charges when they are running—for example, from the time you launch a DB instance until you terminate it.
- **Database characteristics:** The physical capacity of the database you choose will affect how much you are charged. Database characteristics vary depending on the database engine, size, and memory class.
- **Database purchase type:** When you use On-Demand DB Instances, you pay for compute capacity for each hour your DB Instance runs, with no required minimum commitments. With Reserved DB Instances, you can make a low, one-time, upfront payment for each DB Instance you wish to reserve for a 1- or 3-year term.
- **Number of database instances:** With Amazon RDS, you can provision multiple DB instances to handle peak loads.
- **Provisioned storage:** There is no additional charge for backup storage of up to 100 percent of your provisioned database storage for an active DB Instance. After the DB Instance is terminated, backup storage is billed per GB per month.
- **Additional storage:** The amount of backup storage in addition to the provisioned storage amount is billed per GB per month.
- **Requests:** The number of input and output requests to the database.
- **Deployment type:** You can deploy your DB Instance to a single Availability Zone (analogous to a standalone data center) or multiple Availability Zones (analogous to a secondary data center for enhanced availability and durability). Storage and I/O charges vary, depending on the number of Availability Zones you deploy to.
- **Data transfer:** Inbound data transfer is free, and outbound data transfer costs are tiered.

Depending on your application's needs, it's possible to optimize your costs for Amazon RDS database instances by purchasing reserved Amazon RDS database instances. To purchase Reserved Instances, you make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly usage charge for that instance.

Get more information about [Amazon RDS pricing](#).

Amazon DynamoDB

[Amazon DynamoDB](#) is a fast and flexible [NoSQL database](#) service for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed cloud database and supports both document and key-value store models. Its flexible data model, reliable performance, and automatic scaling of throughput capacity make it a great fit for mobile, web, gaming, ad tech, IoT, and many other applications.

Amazon DynamoDB pricing at a glance

Unlike traditional NoSQL deployments that ask you to think about memory, CPU, and other system resources that could affect your throughput, DynamoDB simply asks you to specify the target utilization rate and minimum to maximum capacity that you want for your table. DynamoDB handles the provisioning of resources to achieve your target utilization of read and write capacity, then auto-scales your capacity based on usage. Optionally, you can directly specify read and write capacity if you prefer to manually manage table throughput.

The following table summarizes key DynamoDB pricing concepts.

Resource Type	Details	Monthly Price
Provisioned throughput (write)	One write capacity unit (WCU) provides up to one write per second, enough for 2.5 million writes per month.	As low as \$0.47 per WCU
Provisioned throughput (read)	One read capacity unit (RCU) provides up to two reads per second, enough for 5.2 million reads per month.	As low as \$0.09 per RCU
Indexed data storage	DynamoDB charges an hourly rate per GB of disk space that your table consumes.	As low as \$0.25 per GB

Indexed data storage

Amazon DynamoDB is an SSD-backed indexed data store. The amount of disk space your data consumes will exceed the raw size of the data you have uploaded. Amazon DynamoDB measures the size of your billable data by adding up the raw byte size of the data you upload, plus a per-item storage overhead of 100 bytes to account for indexing. You do not need to “provision” storage with Amazon DynamoDB. You are simply billed for what you use, as described above.

Data transfer

There is no additional charge for data transferred between Amazon DynamoDB and other Amazon Web Services within the same Region. Data transferred across Regions (e.g., between Amazon DynamoDB in the US East (Northern Virginia) Region and Amazon EC2 in the EU (Ireland) Region) will be charged on both sides of the transfer.

Global Tables

[Global Tables](#) builds on DynamoDB's global footprint to provide you with a fully managed, multiregion, and multimaster database that provides fast local read and write performance for massively scaled, global applications. Global Tables replicates your Amazon DynamoDB tables automatically across your choice of AWS regions.

When using Global Tables, you are charged based on the resources associated with each replica table. Write capacity for Global Tables is represented by replicated write capacity units (rWCUs) instead of standard write capacity units (WCUs). For more details on Global Tables, please refer to the [DynamoDB Developer Guide](#).

There is no minimum fee to use DynamoDB. Estimate your monthly bill using the [Simple Monthly Calculator](#).

Learn more about pricing for additional DynamoDB services at the [Amazon DynamoDB pricing page](#).

Amazon CloudFront

[Amazon CloudFront](#) is a global content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to your viewers with low latency and high transfer speeds.

Amazon CloudFront pricing

Amazon CloudFront charges are based on the data transfers and requests used to deliver content to your customers. There are no upfront payments or fixed platform fees, no long-term commitments, no premiums for dynamic content, and no requirements for professional services to get started. There is no charge for data transferred from AWS services such as Amazon S3 or Elastic Load Balancing. And, best of all, you can get started with CloudFront for free.

When you begin to estimate the cost of Amazon CloudFront, consider the following:

- **Traffic distribution:** Data transfer and request pricing varies across geographic regions, and pricing is based on the edge location through which your content is served.
- **Requests:** The number and type of requests (HTTP or HTTPS) made and the geographic region in which the requests are made.
- **Data transfer out:** The amount of data transferred out of your Amazon CloudFront edge locations.

Get more information about [Amazon CloudFront pricing](#).

Optimizing costs with reservations

While the cloud is well suited for variable workloads and rapid deployment, many cloud-based workloads display a more predictable pattern. For these stable applications, organizations can achieve significant cost savings by using Reserved Instances (RIs) and other reservation models for compute and data services. RIs allow you to commit to usage parameters at the time of purchase to achieve a lower hourly rate. For example, using Amazon EC2 Reserved Instances, you can save up to 75 percent over equivalent capacity purchased in an On-Demand model.

Most reservations can be purchased using a range of payment terms including No Upfront, Partial Upfront, or All Upfront. When you buy reservations, the larger the upfront payment, the greater the discount. To maximize your savings, you can pay all up front and receive the largest discount. Partial Upfront RIs offer lower discounts but give you the option to spend less up front. Lastly, you can choose to spend nothing up front and receive a smaller discount, allowing you to free up capital to spend in other projects.

For more information about reservation-based pricing, see the [AWS Cost Optimization page](#).

Amazon EC2 Reserved Instances

You can use [Amazon EC2 Reserved Instances](#) to reserve capacity and receive a discount on your instance usage compared to running On-Demand instances. The discounted usage price is reserved for the duration of your contract, allowing you to predict compute costs over the term of the Reserved Instance. When you want to use the capacity you reserved, you launch an EC2 instance with the same configuration as the reserved capacity that you purchased. AWS will automatically apply the usage price that is associated with your capacity reservation. You are charged that price for your EC2 instance for every hour in the term for as long as you own the Reserved Instance.

Reserved Instance Marketplace

The [Reserved Instance Marketplace](#) allows other AWS customers to list their Reserved Instances for sale. Third-party Reserved Instances are often listed at lower prices and shorter terms. These Reserved Instances are no different from Reserved Instances purchased directly from AWS.

Amazon EC2 savings with Reserved Instances

The following table shows an example of potential savings for Amazon EC2 Reserved Instances for a standard 1- and 3-year term running a single m5.large Linux instance in the US East (N. Virginia) Region.

Standard 1-Year Term

Payment Option	Upfront	Monthly*	Effective Hourly	Savings over On-Demand	On-Demand Hourly
No Upfront	\$0	\$44.53	\$0.061	36%	\$0.096 per Hour
Partial Upfront	\$256	\$21.17	\$0.058	39%	
All Upfront	\$501	\$0	\$0.057	40%	

Standard 3-Year Term

Payment Option	Upfront	Monthly*	Effective Hourly	Savings over On-Demand
No Upfront	\$0	\$30.66	\$0.042	56%
Partial Upfront	\$515	\$14.60	\$0.040	59%
All Upfront	\$968	\$0	\$0.037	62%

As shown the first table, choosing a one-year All Upfront Reserved Instance will save you 40% over On-Demand pricing. The second table shows that three-year RIs savings can increase to 62%.

Amazon DynamoDB Reserved Capacity

If you can predict your need for Amazon DynamoDB read-and-write throughput, Reserved Capacity offers significant savings over the normal price of DynamoDB provisioned throughput capacity. You pay a one-time upfront fee and commit to paying for a minimum usage level at specific hourly rates for the duration of the Reserved Capacity term. Any throughput you provision in excess of your Reserved Capacity is billed at standard rates for provisioned throughput.

Amazon ElastiCache Reserved Nodes

[Amazon ElastiCache Reserved Nodes](#) give you the option to make a low, one-time payment for each cache node you want to reserve and, in turn, receive a significant discount on the hourly charge for that node. Amazon ElastiCache provides three ElastiCache Reserved Node types (Light, Medium, and Heavy Utilization Reserved Nodes) that enable you to balance the amount you pay up front with your effective hourly price.

Amazon RDS RIs

Like Amazon EC2 RIs, Amazon RDS RIs can be purchased using No Upfront, Partial Upfront, or All Upfront terms. All Reserved Instance types are available for Aurora, MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server database engines.

Amazon Redshift Reserved Nodes

If you intend to keep an [Amazon Redshift](#) cluster running continuously for a prolonged period, you should consider purchasing reserved-node offerings. These offerings provide significant savings over on-demand pricing, but they require you to reserve compute nodes and commit to paying for those nodes for either a 1- or 3-year duration.

Conclusion

While 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. Projecting costs for a use case, such as web application hosting, can be challenging, because a solution typically uses multiple features across multiple AWS products, which in turn 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. To help you understand how AWS pricing works in the context of real-world solutions, see a cost-calculation example later in this paper.

You can use the [AWS Simple Monthly 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](#).

Resources

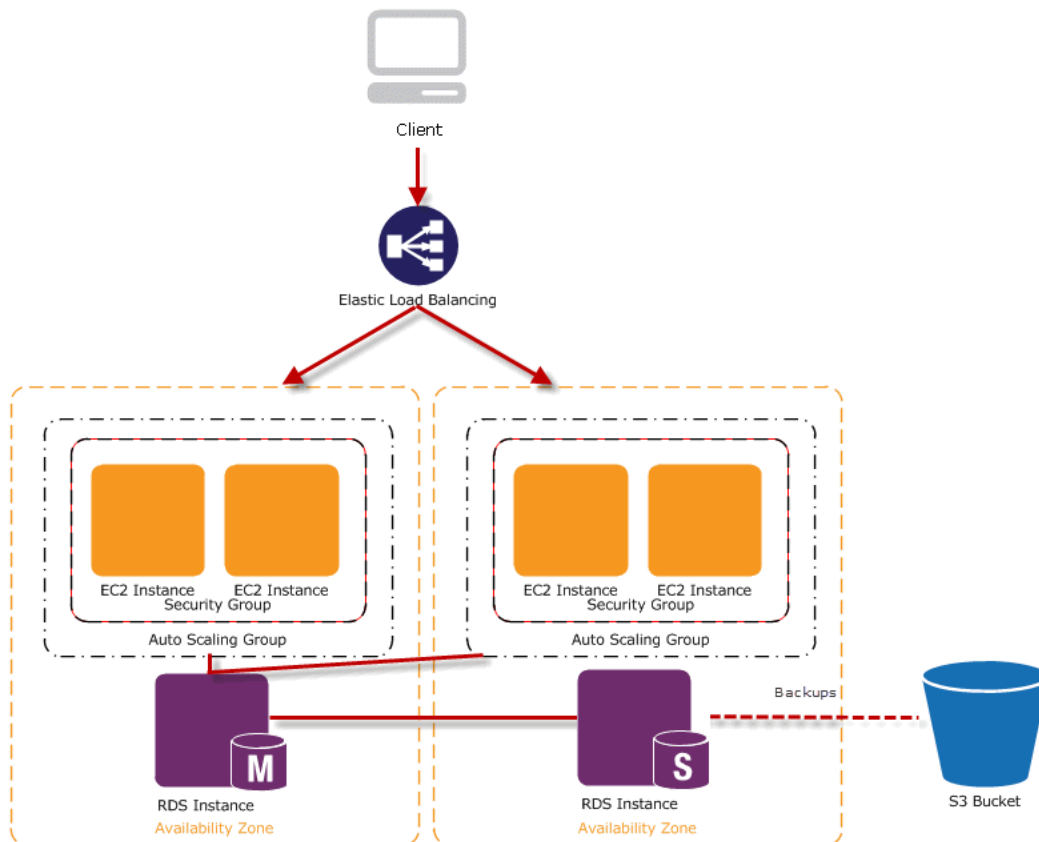
Resource	Description
AWS Simple Monthly Calculator	The AWS Simple Monthly Calculator helps you estimate your monthly bill.
AWS Architecture Center	The AWS Architecture Center provides you with the necessary guidance and best practices to build highly scalable and reliable applications in the AWS Cloud.
Cloud Economics Center	The Cloud Economics Center provides access to information, tools, and resources to compare the costs of Amazon Web Services with IT infrastructure alternatives.
Account	View your current charges and account activity, itemized by service and by usage type. Previous months' billing statements are also available.
AWS Cost and Usage Reports	Reports are available to download for each service. Specifying usage types, timeframe, service operations, and more can customize reports.

Cost calculation example

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

Architecture

Classic Load Balancing balances traffic the Amazon EC2 Instances in an Auto Scaling group, which adds or subtracts Amazon EC2 Instances to match 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, Auto Scaling, and one Amazon RDS database instance across multiple Availability Zones.



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. 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.

Amazon EC2 cost breakdown

The following table shows the characteristics for Amazon EC2 we have identified for this dynamic site in the US-East region.

Characteristic	Estimated Usage	Description
Utilization	100%	The server runs 24 hours a day, 7 days per week.
Machine characteristics	m4.large instance	The instance includes 8 GB of memory, 2 vCPU
Storage	Amazon EBS gp2	We will use 1 EBS volume per instance with 4GB of storage per volume.
Data backup	Daily Amazon EBS snapshots	10% incremental change per day
Data transfer	Data In: 5 GB/day Data Out: 50 GB/day	There are approximately 1,000,000 hits per day. Each response transfers out about 50 KB, and each request transfers in about 5 KB.
Instance scale	4	On average in each day, there are 4 instances running.
Load Balancing	55 GB/day	Classic Load Balancing is used 24 hours per day, 7 days per week. It processes a total of 55 GB/day (data in + data out).

The total cost for one month is the sum of the cost of the running instances, Amazon EBS storage and snapshots, load balancing, and data transfer out, minus the Free Tier discount. We calculated the total cost using the AWS Simple Monthly Calculator.

Service	Monthly price
Amazon EC2 Compute	\$585.60
Amazon EBS Volumes	\$1.60
Amazon EBS snapshots	\$2.08
Classic Load Balancers	\$18.30
Data Processed by Classic Load Balancers:	\$13.42
AWS Data Transfer Out	\$137.16
Free Tier Discount	(\$21.33)
Total Monthly Payment	\$736.83