Amazon Athena pricing

Why Athena?

Amazon Athena is a serverless, interactive analytics service built on open-source frameworks that enables you to analyze petabytes of data where it lives. With Athena, you can use SQL or Apache Spark and there is no infrastructure to set up or manage. Pricing is simple: you pay based on data processed or compute used.

To get started, you create a workgroup that will allow you to specify your query engine, your working directory in Amazon Simple Storage Service (S3) to hold the results of your execution, AWS Identity and Access Management (IAM) roles (if needed), and your resource tags. You can use workgroups to separate users, teams, applications, or workloads; set limits on the amount of data that each query or the entire workgroup can process; and track costs. Based on the workgroup that you create, you can either (a) run SQL queries and get based on data scanned or compute used or (b) run Apache Spark Python code and get charged an hourly rate for executing your code.

AWS pricing calculator

Calculate your Amazon Athena and architecture cost in a single estimate.

Create your custom estimate now

SQL queries

Region:

SQL queries with Provisioned Capacity

Region:

Apache Spark

Region:

Additional costs

Athena queries data directly from Amazon S3. There are no additional storage charges for querying your data with Athena. You are charged standard S3 rates for storage, requests, and data transfer. By default, query results are stored in an S3 bucket of your choice and are also billed at standard S3 rates.

- You are billed by S3 when your workloads read, store, and transfer data. This includes successful and unsuccessful queries. By default, SQL query results and Spark calculation results are stored in an S3 bucket of your choice and billed at standard S3 rates. See Amazon S3 pricing for more information.
- If you use the AWS Glue Data Catalog with Athena, you are charged standard Data Catalog rates. For details, visit the AWS Glue pricing page.
- SQL queries on federated data sources (data not stored on S3) are billed per terabyte (TB) scanned by Athena aggregated across data sources, rounded up to the nearest megabyte with a 10 megabyte minimum per query, unless Provisioned Capacity is used. Such queries also invoke AWS Lambda functions in your account, and you are charged for Lambda use at standard rates. Lambda functions invoked by federated queries are subject to Lambda's free tier. Visit the Lambda pricing page for details.

Pricing examples

Example 1 – SQL query

Consider a table with 4 equally sized columns, stored as an uncompressed text file with a total size of 3 TB on Amazon S3. Running a query to get data from a single column of the table requires Amazon Athena to scan the entire file because text formats can't be split.

- This query would cost: $15. (Price for 3 TB scanned is 3 * $5/TB = $15.)

If you compress your file using GZIP, you might see 3:1 compression gains. In this case, you would have a compressed file with a size of 1 TB. The same query on this file would cost $5. Athena has to scan the entire file again, but because it's three times smaller in size, you pay one-third of what you did before. If you compress your file and also convert it to a columnar format like Apache Parquet, achieving 3:1 compression, you would still end up with 1 TB of data on S3. But, in this case, because Parquet is columnar, Athena can read only the column that is relevant for the query being run. Because the query in question only references a single column, Athena reads only that column and can avoid reading three-fourths of the file. Since Athena only reads one-fourth of the file, it scans just 0.25TB of data from S3.

- This query would cost: $1.25. There is a 3x savings from compression and 4x savings for reading only one column.
  (File size = 3TB/3 = 1 TB. Data scanned when reading a single column = 1TB/4 = 0.25 TB. Price for 0.25 TB = 0.25 * $5/TB = $1.25)

Example 2 – SQL queries with Provisioned Capacity

Suppose your team supports a web application that provides self-service analytics to users who submit queries during business hours and expect their queries to complete in a predictable amount of time. Last week, application users submitted a total of 10,000 queries which scanned 500 TB of data. You want to use Provisioned Capacity to help you maintain a consistent user experience as the number of users grows. From analysis of your queries, you determine that 96 DPU are sufficient for your current workload.

- For one business day, the cost to support this workload with Provisioned Capacity is calculated as 96 DPU * $0.30 per DPU Hour * 12 hours per day = $345.60 USD.

One morning, you learn that a new set of application users has completed onboarding and, as a result, you expect query volume to be 2x higher than it was the day before. You want to ensure users have similar performance as yesterday, but don't expect all users to submit queries at the same time. Two hours into the day, you increase capacity by 50% to 144 DPU.

- The cost for today's workload is equal to the cost of 96 DPU for 2 hours plus 144 DPU for 10 hours, or 96 DPU * $0.30 per DPU Hour * 2 hours + 144 DPU * $0.30 per DPU Hour * 10 hours = $489.60 USD.

Example 3 – Apache Spark application

Consider using a notebook in the Athena console for pulling sales figures for the previous quarter and graphing them to create a report. You start a session using a notebook. Your session lasts for 1 hour and submits 6 calculations as part of the session. Each calculation takes 20 1-DPU worker nodes to run and lasts for 1 minute.

- Worker DPU-hours = Number of calculations * DPUs used per calculations * execution time of calculation = 6 calculations * 20 DPUs per calculation * (1/60) hours per calculation = 2.0 DPU-hours
- Driver DPU-hours = DPUs used per session * session time = 1 DPUs per session * 1 hours per session = 1.0 DPU-hours
- Total DPU-hours = Worker DPU-hours + Driver DPU-hours = 2.0 DPU-hours + 1.0 DPU-hours = 3.0 DPU-hours
- Spark application charges = $0.35 per DPU-hour * 3.0 DPU-hours = $1.05

Note: S3 will charge you separately to store and read your data and the results of your execution.

Additional pricing resources

AWS Pricing Calculator
Easily calculate your monthly costs with AWS

Get pricing assistance
Contact AWS specialists to get a personalized quote

Visit the getting started page
Learn how to get started with Amazon Athena.
Learn more »

Start building in the console
Get started building with Amazon Athena in the AWS Management Console.
Sign in »

Contact an expert
Explore the Amazon Athena support options.
Contact us »

Get Started with AWS
Learn how to start using AWS in minutes

AWS Free Tier
Gain free, hands-on experience with AWS for 12 months

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