AWS Lambda pricing

Overview

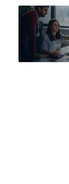

With Lambda, you can run code for virtually any type of application or backend service - all with zero administration. Just upload your code as a function and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services or call it directly from any web or mobile app. You pay only for the number of requests for your functions and the duration it takes for your code to execute.

Lambda counts a request each time it starts executing in response to an event notification or an invoke call, including test invokes from the console. Duration is calculated from the time your code begins executing until it returns or otherwise terminates, rounded up to the nearest 1 ms. The price depends on the amount of memory you allocate to your function. In the AWS GovCloud (US) Regions, you pay a flat fee of $0.20 per 1 million requests and $0.0000166667 for every GB-second used.

You can use your Lambda functions to process built-in AWS or 3rd-party cloud services or your own applications. For example, you can use Lambda to process data events from Amazon DynamoDB, Amazon S3, Amazon Kinesis, Amazon SNS, Alexa, or any Web, mobile, or IoT application. Lambda automatically scales your application by running code in response to each trigger. Your code runs in parallel and processes each trigger individually, scaling precisely with the size of the workload.

Lambda's free usage tier includes one million free requests per month and 400,000 GB-seconds of compute time per month. You can use this free tier to build a serverless application using any combination of Lambda functions and integrated AWS services. Additionally, the Lambda free usage tier includes the first 1 million requests and 400,000 GB-seconds of compute time per month for functions running in Lambda@Edge. Unused free tier minutes and requests do not carry forward to subsequent months. Learn more about [AWS Free Tier](#) and [Lambda@Edge](#).

To learn more about Lambda, see the [Lambda FAQs](#) and [documentation](#).

AWS Pricing Calculator

Calculate your AWS estimate and architecture cost for a single workload.

[Create your custom estimate now](#)

AWS Lambda Pricing

Region:

Europe (London)

Architecture	Duration	Requests
x86 Price		
First 6 Billion GB-seconds / month	$0.0000133667 for every GB-second	$0.20 per 1M requests
Next 9 Billion GB-seconds / month	$0.0000113 for every GB-second	$0.20 per 1M requests
Over 15 Billion GB-seconds / month	$0.0000106334 for every GB-second	$0.20 per 1M requests
Arm Price		
First 5 Billion GB-seconds / month	$0.0000133334 for every GB-second	$0.20 per 1M requests
Next 9.5 Billion GB-seconds / month	$0.0000113000 for every GB-second	$0.20 per 1M requests
Over 14.5 Billion GB-seconds / month	$0.0000106334 for every GB-second	$0.20 per 1M requests

Asynchronous event (including events from S3, SNS, EventBridge, StepFunctions, CloudWatch Logs): You are charged for 1 request per each asynchronous event that Lambda processes. Lambda counts a request each time it starts executing in response to an event notification or an invoke call.

Duration: Duration depends on the amount of memory you allocate to your function. You are charged for the amount of memory in use from the nearest 1 ms. Duration is calculated from the time your code begins executing until it returns or otherwise terminates, rounded up to the nearest 1 ms. The price depends on the amount of memory you allocate to your function. The Lambda free usage tier includes 400,000 GB-seconds of compute time per month.

Asynchronous event (including events from S3, SNS, EventBridge, StepFunctions, CloudWatch Logs): You are charged for 1 request per each asynchronous event that Lambda processes. Lambda counts a request each time it starts executing in response to an event notification or an invoke call.

Duration: Duration depends on the amount of memory you allocate to your function. You are charged for the amount of memory in use from the nearest 1 ms. Duration is calculated from the time your code begins executing until it returns or otherwise terminates, rounded up to the nearest 1 ms. The price depends on the amount of memory you allocate to your function. In the GovCloud Regions, you pay a flat fee of $0.20 per 1 million requests and $0.0000166667 for every GB-second used. The Lambda free usage tier includes 400,000 GB-seconds of compute time per month.

Example 1: You allocate 512 MB of memory to your function and execute it in your region for 3 million times per month for 1 second per execution. Your monthly compute usage is 3 million * 1 second * 0.5 GB = 1,536,000 GB-seconds. Your monthly compute charges are 1,536,000 GB-seconds * $0.0000166667 = $25.60. Your monthly request charges are 3 million * $0.20 per 1 million = $0.60. Your total monthly charges are $25.60 + $0.60 = $26.20.

Example 2: You allocate 1 GB of memory to your function and execute it in your region for 3 million times per month for 1 second per execution. Your monthly compute usage is 3 million * 1 second * 1 GB = 3,000,000 GB-seconds. Your monthly compute charges are 3,000,000 GB-seconds * $0.0000166667 = $50.00. Your monthly request charges are 3 million * $0.20 per 1 million = $0.60. Your total monthly charges are $50.00 + $0.60 = $50.60.

Example 3: You allocate 1 GB of memory to your function and execute it in your region for 3 million times per month for 1 minute per execution. Your monthly compute usage is 3 million * 60 seconds * 1 GB = 180,000,000 GB-seconds. Your monthly compute charges are 180,000,000 GB-seconds * $0.0000166667 = $3,000.00. Your monthly request charges are 3 million * $0.20 per 1 million = $0.60. Your total monthly charges are $3,000.00 + $0.60 = $3,000.60.

Asynchronous event (including events from S3, SNS, EventBridge, StepFunctions, CloudWatch Logs): You are charged for 1 request per each asynchronous event that Lambda processes. Lambda counts a request each time it starts executing in response to an event notification or an invoke call.

Duration: Duration depends on the amount of memory you allocate to your function. You are charged for the amount of memory in use from the nearest 1 ms. Duration is calculated from the time your code begins executing until it returns or otherwise terminates, rounded up to the nearest 1 ms. The price depends on the amount of memory you allocate to your function. In the AWS GovCloud (US)