Amazon Lex pricing

Feedback

Amazon Lex is a service for building conversational interfaces into any application using voice and text. With Amazon Lex, you pay only for what you use. There is no upfront commitment or minimum fee. Amazon Lex bots are designed for a request and response interaction or a continuous streaming conversation. With the request and response interaction, each user input (voice or text) is processed as a separate API call. In a streaming conversation, all user inputs across multiple turns are processed in one streaming API call. Pricing for Amazon Lex is based upon these two interaction models.

With streaming conversation, the bot continuously listens and can be designed to respond proactively. For example, you can configure the bot to keep a conversation going when the user needs more time to respond by sending periodic messages such as "Take your time. Let me know once you are ready." The request and response model is a different user experience where a user input is required as an initiator.

Starting July 15, 2025, new AWS customers will receive up to $200 in AWS Free Tier credits, which can be applied towards eligible AWS services, including Amazon Lex. At account sign-up, you can choose between a free plan and a paid plan. The free plan will be available for 6 months after account creation. If you upgrade to a paid plan, any remaining Free Tier credit balance will automatically apply to your AWS bills. All Free Tier credits must be used within 12 months of your account creation date. To learn more about the AWS Free Tier program, refer to AWS Free Tier website and AWS Free Tier documentation.

Request and response interaction pricing

Region:

Europe (London)

With the request and response interaction, each user input is processed as a separate API call. You are charged based on the number of speech or text API requests processed by your bot at $0.004 per speech request and $0.00075 per text request. For example, the cost for 1,000 speech requests would be $4.00, and 1,000 text requests would cost $0.75. The speech and text requests are added up at the end of the month to generate your monthly charges.

Key terms

Speech request: Each speech input from the user is counted as a speech request.
Text request: Each text input from the user is counted as a text request.

Pricing Example

Consider a bot that processes 8,000 speech requests and 2,000 text requests in one month.

Input requests	Cost per request	Number of requests	Total
8,000 speech requests	$0.004	8,000 requests	$32.00
2,000 text requests	$0.00075	2,000 requests	$1.50
Total Amazon Lex charges for the month			**$33.50**

Streaming conversation pricing

Region:

Europe (London)

In a streaming conversation, all user inputs across multiple turns are processed in one streaming API call. The bot continuously listens and can be designed to respond proactively. For example, when capturing information that user may not have readily available (e.g., credit card number), you can design the bot to respond with periodic messages such as "Let me know when you are ready with the card details." Streaming speech input is billed at $0.0065 per 15-second interval for the duration of the open streaming connection, while Amazon Lex is actively listening for audio input. This includes active speech, pause detection, silence thresholds, processing time, and any wait states. Billing continues as long as your connection remains open and Lex is ready to receive input. Bot responses are not counted towards the usage. For example, consider user input that lasts for 52 seconds including speech and silence. It will be rounded off to 60 seconds and counted as four 15-second speech intervals at $0.0065 per 15-second interval for a total cost of $0.026. For streaming text input, you are charged based on the number of text requests processed by your bot at $0.002 per request. The speech intervals and text requests are added up at the end of the month to generate your monthly charges.

Key terms

Speech intervals: Every 15 seconds of input (including speech and silence) is counted as one interval; any input is rounded up to the nearest 15-second interval.
Text requests: One text user input is counted as a text request.

Pricing Example

Consider a bot that processes 8,000 speech intervals and 2,000 text requests in one month.

Input requests	Cost per unit	Number of units	Total
8,000 speech intervals*	$0.0065	8,000 intervals	$52.00
2,000 text requests	$0.0020	2,000 requests	$4.00
Total Amazon Lex charges for the month			$56.00

*Every 15 seconds of input (including silence) is counted as one interval; any input is rounded up to the nearest 15-second interval.

Automated chatbot designer pricing

Region:

Europe (London)

The Amazon Lex Automated chatbot designer helps you automatically create chatbots using conversation transcripts within hours. It uses natural language understanding (NLU) and deep learning techniques to analyze conversation transcripts and surfaces a chatbot design that includes user intents (e.g., "file a claim" for an insurance bot), sample phrases (e.g., "I'd like to file a claim") associated with those intents, and a list of all the information (e.g., claim type, policy ID) required to fulfill them. You can iterate on this design and change chatbot prompts and responses instead of starting from scratch.

You will be charged a training fee of $0.50 per minute for the time it takes to analyze the transcripts and discover intents. The amount of time it takes to complete training depends on a number of factors including the volume of transcripts and the complexity of the conversations. Typically, 500 lines of transcripts are analyzed every minute. The training time is rounded up to the nearest 1-minute interval. For example, if it takes 15 minutes and 30 seconds to complete the training and surface a chatbot design, the training time will be rounded off to 16 minutes. At $0.50 per minute, your total cost for this time will be $8.00.

Key terms

Transcripts: A text file listing conversations between the caller and the agent. Each interaction in the conversation can span multiple lines.
Training time: Time required to analyze transcripts and identify intents.

Pricing examples

Example 1

You are a regional credit union and operate a contact center to help customers with queries and transactions related to their bank accounts. You want create a bot to augment your contact center operations and improve efficiencies. You select the conversation transcripts from customer calls handled by your high performing agents as an input to the automated chatbot designer to create a high-quality bot design. Let's assume that the calls handled by the top tier agents over a month generate 180K lines of transcripts (assume your contact center has 100 agents, you pick the top 10 high performing agents each answering 20 calls/day; assuming 30 lines of transcript per call, you will have 6K lines daily or 180K lines monthly). The automated chatbot designer takes about 5 hours (or 300 minutes) to analyze the conversation transcripts and surface the design. You are charged for the 300 minutes of training time at $0.50 per minute, leading to total training charges of $150.00 for a month for the 180K lines of transcripts.

Total training time	Charges per minute	Total charges
300 minutes	$0.50	$150.00

Example 2

You are an insurance company with a contact center providing customer support to auto, home, and life insurance policy holders. You want to automate auto insurance conversations with a bot that can help customers with transactions such as making premium payments and filing claims. You use the conversation transcripts from calls with a high customer satisfaction score (CSAT) to ensure high-quality input to the automated chatbot designer. Let's assume you have 600K lines of transcripts from customer conversations over a week with a high customer satisfaction score (assume a contact center size of 1,500 agents answering a total of 200K calls per week, assume 20K of these calls receive 90% or higher CSAT; with 30 lines per transcript, this results in 600K total lines). The automated chatbot designer takes about 17 hours (or 1,000 minutes) to analyze the conversation transcripts and surface the design. You are charged for 1,080 minutes of training time at $0.50 per minute, leading to total training charges of $540 for the 600K lines of conversation transcripts.

Total training time	Charges per minute	Total charges
1,000 minutes	$0.50	$500.00

Additional pricing resources

AWS Pricing Calculator
Easily calculate your monthly costs with AWS.

Get pricing assistance
Contact AWS specialists to get a personalized quote.

Next steps

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