



Developer Guide

Amazon Lex V1



Amazon Lex V1: Developer Guide

Copyright © 2025 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

.....	viii
What Is Amazon Lex?	1
Are You a First-time User of Amazon Lex?	2
How It Works	4
Supported Languages	6
Supported Languages and Locales	7
Languages and Locales Supported by Amazon Lex Features	7
Programming Model	8
Model Building API Operations	8
Runtime API Operations	9
Lambda Functions As Code Hooks	11
Managing Messages	13
Types of Messages	14
Contexts for Configuring Messages	15
Supported Message Formats	20
Message Groups	20
Response Cards	22
Managing Conversation Context	26
Setting Intent Context	27
Using Default Slot Values	29
Setting Session Attributes	30
Setting Request Attributes	32
Setting the Session Timeout	35
Sharing Information Between Intents	36
Setting Complex Attributes	36
Using Confidence Scores	38
Session Management	40
Conversation Logs	41
IAM Policies for Conversation Logs	42
Configuring Conversation Logs	45
Encrypting Conversation Logs	49
Viewing Text Logs in Amazon CloudWatch Logs	50
Accessing Audio Logs in Amazon S3	54
Monitoring Conversation Log Status with CloudWatch Metrics	55

Managing Sessions	55
Switching Intents	57
Resuming a Prior Intent	58
Starting a New Session	59
Validating Slot Values	59
Deployment Options	59
Built-in Intents and Slot Types	60
Built-in Intents	60
Built-in Slot Types	78
Custom Slot Types	89
Slot Obfuscation	91
Sentiment Analysis	92
Tagging Resources	93
Tagging Your Resources	94
Tag Restrictions	94
Tagging Resources (Console)	95
Tagging Resources (AWS CLI)	97
Getting Started	99
Step 1: Set Up an Account	99
Sign Up for AWS	99
Create a user	100
Next Step	101
Step 2: Set Up the AWS CLI	101
.....	102
Step 3: Getting Started (Console)	102
Exercise 1: Create a Bot Using a Blueprint	103
Exercise 2: Create a Custom Bot	140
Exercise 3: Publish a Version and Create an Alias	156
Step 4: Getting Started (AWS CLI)	157
Exercise 1: Create a Bot	158
Exercise 2: Add a New Utterance	176
Exercise 3: Add a Lambda Function	181
Exercise 4: Publish a Version	185
Exercise 5: Create an Alias	192
Exercise 6: Clean Up	193
Versioning and Aliases	195

Versioning	195
The \$LATEST Version	195
Publishing an Amazon Lex Resource Version	196
Updating an Amazon Lex Resource	197
Deleting an Amazon Lex Resource or Version	197
Aliases	198
Using Lambda Functions	200
Lambda Function Input Event and Response Format	200
Input Event Format	200
Response Format	208
Amazon Lex and AWS Lambda Blueprints	215
Updating a Blueprint for a Specific Locale	216
Deploying Bots	217
Deploying an Amazon Lex Bot on a Messaging Platform	217
Integrating with Facebook	220
Integrating with Kik	223
Integrating with Slack	227
Integrating with Twilio SMS	233
Deploying an Amazon Lex Bot in Mobile Applications	236
Importing and Exporting	237
Exporting and Importing in Amazon Lex Format	237
Exporting in Amazon Lex Format	238
Importing in Amazon Lex Format	239
JSON Format for Importing and Exporting	241
Exporting to an Alexa Skill	244
Bot Examples	246
Schedule Appointment	246
Overview of the Bot Blueprint (ScheduleAppointment)	249
Overview of the Lambda Function Blueprint (lex-make-appointment-python)	250
Step 1: Create an Amazon Lex Bot	251
Step 2: Create a Lambda Function	253
Step 3: Update the Intent: Configure a Code Hook	254
Step 4: Deploy the Bot on the Facebook Messenger Platform	255
Details of Information Flow	256
Book Trip	274
Step 1: Blueprint Review	275

Step 2: Create an Amazon Lex Bot	278
Step 3: Create a Lambda function	281
Step 4: Add the Lambda Function as a Code Hook	282
Details of the Information Flow	286
Example: Using a Response Card	306
Updating Utterances	310
Integrating with a Web site	312
Call Center Agent Assistant	312
Step 1: Create an Amazon Kendra Index	314
Step 2: Create an Amazon Lex Bot	314
Step 3: Add a Custom and Built-in Intent	315
Step 4: Set up Amazon Cognito	316
Step 5: Deploy Your Bot as a Web Application	318
Step 6: Use the Bot	318
Migrating a bot	322
Migrating a bot (Console)	322
Migrating a Lambda function	323
Migration messages	324
Built-in intent	324
Built-in slot type	324
Conversation logs	324
Message groups	325
Prompts and phrases	325
Other Amazon Lex V1 features	326
Migrating a Lambda function	326
List of updated fields	328
Security	336
Data Protection	337
Encryption at Rest	337
Encryption in Transit	339
Key Management	339
Identity and Access Management	339
Audience	339
Authenticating with identities	340
Managing access using policies	343
How Amazon Lex works with IAM	346

Identity-based policy examples	357
AWS managed policies for Amazon Lex	363
Using Service-Linked Roles	372
Troubleshooting	374
Monitoring	375
Monitoring Amazon Lex with CloudWatch	376
Logging Amazon Lex API Calls with AWS CloudTrail	388
Compliance Validation	393
Resilience	394
Infrastructure Security	394
Guidelines and Quotas	395
Supported Regions	395
General Guidelines	395
Quotas	398
Runtime Service Quotas	399
Model Building Quotas	401
API Reference	406
Actions	406
Amazon Lex Model Building Service	408
Amazon Lex Runtime Service	620
Data Types	664
Amazon Lex Model Building Service	665
Amazon Lex Runtime Service	723
Document History	742
AWS Glossary	749

End of support notice: On September 15, 2025, AWS will discontinue support for Amazon Lex V1. After September 15, 2025, you will no longer be able to access the Amazon Lex V1 console or Amazon Lex V1 resources. If you are using Amazon Lex V2, refer to the [Amazon Lex V2 guide](#) instead. .

What Is Amazon Lex?

Amazon Lex is an AWS service for building conversational interfaces for applications using voice and text. With Amazon Lex, the same conversational engine that powers Amazon Alexa is now available to any developer, enabling you to build sophisticated, natural language chatbots into your new and existing applications. Amazon Lex provides the deep functionality and flexibility of natural language understanding (NLU) and automatic speech recognition (ASR) so you can build highly engaging user experiences with lifelike, conversational interactions, and create new categories of products.

Amazon Lex enables any developer to build conversational chatbots quickly. With Amazon Lex, no deep learning expertise is necessary—to create a bot, you just specify the basic conversation flow in the Amazon Lex console. Amazon Lex manages the dialogue and dynamically adjusts the responses in the conversation. Using the console, you can build, test, and publish your text or voice chatbot. You can then add the conversational interfaces to bots on mobile devices, web applications, and chat platforms (for example, Facebook Messenger).

Amazon Lex provides pre-built integration with AWS Lambda, and you can easily integrate with many other services on the AWS platform, including Amazon Cognito, AWS Mobile Hub, Amazon CloudWatch, and Amazon DynamoDB. Integration with Lambda provides bots access to pre-built serverless enterprise connectors to link to data in SaaS applications, such as Salesforce, HubSpot, or Marketo.

Some of the benefits of using Amazon Lex include:

- **Simplicity** – Amazon Lex guides you through using the console to create your own chatbot in minutes. You supply just a few example phrases, and Amazon Lex builds a complete natural language model through which the bot can interact using voice and text to ask questions, get answers, and complete sophisticated tasks.
- **Democratized deep learning technologies** – Powered by the same technology as Alexa, Amazon Lex provides ASR and NLU technologies to create a Speech Language Understanding (SLU) system. Through SLU, Amazon Lex takes natural language speech and text input, understands the intent behind the input, and fulfills the user intent by invoking the appropriate business function.

Speech recognition and natural language understanding are some of the most challenging problems to solve in computer science, requiring sophisticated deep learning algorithms to be trained on massive amounts of data and infrastructure. Amazon Lex puts deep learning technologies within reach of all developers, powered by the same technology as Alexa. Amazon Lex chatbots convert incoming speech to text and understand the user intent to generate an intelligent response, so you can focus on building your bots with differentiated value-add for your customers, to define entirely new categories of products made possible through conversational interfaces.

- **Seamless deployment and scaling** – With Amazon Lex, you can build, test, and deploy your chatbots directly from the Amazon Lex console. Amazon Lex enables you to easily publish your voice or text chatbots for use on mobile devices, web apps, and chat services (for example, Facebook Messenger). Amazon Lex scales automatically so you don't need to worry about provisioning hardware and managing infrastructure to power your bot experience.
- **Built-in integration with the AWS platform** – Amazon Lex has native interoperability with other AWS services, such as Amazon Cognito, AWS Lambda, Amazon CloudWatch, and AWS Mobile Hub. You can take advantage of the power of the AWS platform for security, monitoring, user authentication, business logic, storage, and mobile app development.
- **Cost-effectiveness** – With Amazon Lex, there are no upfront costs or minimum fees. You are charged only for the text or speech requests that are made. The pay-as-you-go pricing and the low cost per request make the service a cost-effective way to build conversational interfaces. With the Amazon Lex free tier, you can easily try Amazon Lex without any initial investment.

Are You a First-time User of Amazon Lex?

If you are a first-time user of Amazon Lex, we recommend that you read the following sections in order:

1. [Getting Started with Amazon Lex](#) – In this section, you set up your account and test Amazon Lex.

2. [API Reference](#) – This section provides additional examples that you can use to explore Amazon Lex.

Amazon Lex: How It Works

Amazon Lex enables you to build applications using a speech or text interface powered by the same technology that powers Amazon Alexa. Following are the typical steps you perform when working with Amazon Lex:

1. Create a bot and configure it with one or more intents that you want to support. Configure the bot so it understands the user's goal (intent), engages in conversation with the user to elicit information, and fulfills the user's intent.
2. Test the bot. You can use the test window client provided by the Amazon Lex console.
3. Publish a version and create an alias.
4. Deploy the bot. You can deploy the bot on platforms such as mobile applications or messaging platforms such as Facebook Messenger.

Before you get started, familiarize yourself with the following Amazon Lex core concepts and terminology:

- **Bot** – A bot performs automated tasks such as ordering a pizza, booking a hotel, ordering flowers, and so on. An Amazon Lex bot is powered by Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU) capabilities. Each bot must have a unique name within your account.

Amazon Lex bots can understand user input provided with text or speech and converse in natural language. You can create Lambda functions and add them as code hooks in your intent configuration to perform user data validation and fulfillment tasks.

- **Intent** – An intent represents an action that the user wants to perform. You create a bot to support one or more related intents. For example, you might create a bot that orders pizza and drinks. For each intent, you provide the following required information:
 - **Intent name**– A descriptive name for the intent. For example, **OrderPizza**. Intent names must be unique within your account.

- **Sample utterances** – How a user might convey the intent. For example, a user might say "Can I order a pizza please" or "I want to order a pizza".
- **How to fulfill the intent** – How you want to fulfill the intent after the user provides the necessary information (for example, place order with a local pizza shop). We recommend that you create a Lambda function to fulfill the intent.

You can optionally configure the intent so Amazon Lex simply returns the information back to the client application to do the necessary fulfillment.

In addition to custom intents such as ordering a pizza, Amazon Lex also provides built-in intents to quickly set up your bot. For more information, see [Built-in Intents and Slot Types](#).

- **Slot** – An intent can require zero or more slots or parameters. You add slots as part of the intent configuration. At runtime, Amazon Lex prompts the user for specific slot values. The user must provide values for all *required* slots before Amazon Lex can fulfill the intent.

For example, the `OrderPizza` intent requires slots such as pizza size, crust type, and number of pizzas. In the intent configuration, you add these slots. For each slot, you provide slot type and a prompt for Amazon Lex to send to the client to elicit data from the user. A user can reply with a slot value that includes additional words, such as "large pizza please" or "let's stick with small." Amazon Lex can still understand the intended slot value.

- **Slot type** – Each slot has a type. You can create your custom slot types or use built-in slot types. Each slot type must have a unique name within your account. For example, you might create and use the following slot types for the `OrderPizza` intent:
 - **Size** – With enumeration values `Small`, `Medium`, and `Large`.
 - **Crust** – With enumeration values `Thick` and `Thin`.

Amazon Lex also provides built-in slot types. For example, `AMAZON.NUMBER` is a built-in slot type that you can use for the number of pizzas ordered. For more information, see [Built-in Intents and Slot Types](#).

For a list of AWS Regions where Amazon Lex is available, see [AWS Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

The following topics provide additional information. We recommend that you review them in order and then explore the [Getting Started with Amazon Lex](#) exercises.

Topics

- [Languages Supported in Amazon Lex](#)
- [Programming Model](#)
- [Managing Messages](#)
- [Managing Conversation Context](#)
- [Using Confidence Scores](#)
- [Conversation Logs](#)
- [Managing Sessions With the Amazon Lex API](#)
- [Bot Deployment Options](#)
- [Built-in Intents and Slot Types](#)
- [Custom Slot Types](#)
- [Slot Obfuscation](#)
- [Sentiment Analysis](#)
- [Tagging Your Amazon Lex Resources](#)

Languages Supported in Amazon Lex

Amazon Lex V1 supports a variety of languages and locales. The languages supported and the features that support them are listed in the following tables.

Amazon Lex V2 supports additional languages, see [Languages Supported in Amazon Lex V2](#)

Supported Languages and Locales

Amazon Lex V1 supports the following languages and locales.

Code	Language and locale
de-DE	German (German)
en-AU	English (Australia)
en-GB	English (UK)
en-IN	English (India)
en-US	English (US)
es-419	Spanish (Latin America)
es-ES	Spanish (Spain)
es-US	Spanish (US)
fr-CA	French (Canada)
fr-FR	French (France)
it-IT	Italian (Italy)
ja-JP	Japanese (Japan)
ko-KR	Korean (Korea)

Languages and Locales Supported by Amazon Lex Features

All Amazon Lex features are supported in all languages and locales except as listed in this table.

Feature	Supported languages and locales
Setting Intent Context	English (US) (en-US)

Programming Model

A *bot* is the primary resource type in Amazon Lex. The other resource types in Amazon Lex are *intent*, *slot type*, *alias*, and *bot channel association*.

You create a bot using the Amazon Lex console or the model building API. The console provides a graphical user interface that you use to build a production-ready bot for your application. If you prefer, you can use the model building API through the AWS CLI or your own custom program to create a bot.

After you create a bot, you deploy it on one of the [supported platforms](#) or integrate it into your own application. When a user interacts with the bot, the client application sends requests to the bot using the Amazon Lex runtime API. For example, when a user says "I want to order pizza," your client sends this input to Amazon Lex using one of the runtime API operations. Users can provide input as speech or text.

You can also create Lambda functions and use them in an intent. Use these Lambda function code hooks to perform runtime activities such as initialization, validation of user input, and intent fulfillment. The following sections provide additional information.

Topics

- [Model Building API Operations](#)
- [Runtime API Operations](#)
- [Lambda Functions As Code Hooks](#)

Model Building API Operations

To programmatically create bots, intents, and slot types, use the model building API operations. You can also use the model building API to manage, update, and delete resources for your bot. The model building API operations include:

- [PutBot](#), [PutBotAlias](#), [PutIntent](#), and [PutSlotType](#) to create and update bots, bot aliases, intents, and slot types, respectively.
- [CreateBotVersion](#), [CreateIntentVersion](#), and [CreateSlotTypeVersion](#) to create and publish versions of your bots, intents, and slot types, respectively.
- [GetBot](#) and [GetBots](#) to get a specific bot or a list of bots that you have created, respectively.

- [GetIntent](#) and [GetIntents](#) to get a specific intent or a list of intents that you have created, respectively.
- [GetSlotType](#) and [GetSlotTypes](#) to get a specific slot type or a list of slot types that you have created, respectively.
- [GetBuiltinIntent](#), [GetBuiltinIntents](#), and [GetBuiltinSlotTypes](#) to get an Amazon Lex built-in intent, a list of Amazon Lex built-in intents, or a list of built-in slot types that you can use in your bot, respectively.
- [GetBotChannelAssociation](#) and [GetBotChannelAssociations](#) to get an association between your bot and a messaging platform or a list of the associations between your bot and messaging platforms, respectively.
- [DeleteBot](#), [DeleteBotAlias](#), [DeleteBotChannelAssociation](#), [DeleteIntent](#), and [DeleteSlotType](#) to remove unneeded resources in your account.

You can use the model building API to create custom tools to manage your Amazon Lex resources. For example, there is a limit of 100 versions each for bots, intents, and slot types. You could use the model building API to build a tool that automatically deletes old versions when your bot nears the limit.

To make sure that only one operation updates a resource at a time, Amazon Lex uses checksums. When you use a Put API operation—[PutBot](#), [PutBotAlias](#), [PutIntent](#), or [PutSlotType](#)—to update a resource, you must pass the current checksum of the resource in the request. If two tools try to update a resource at the same time, they both provide the same current checksum. The first request to reach Amazon Lex matches the current checksum of the resource. By the time that the second request arrives, the checksum is different. The second tool receives a `PreconditionFailedException` exception and the update terminates.

The Get operations—[GetBot](#), [GetIntent](#), and [GetSlotType](#)—are eventually consistent. If you use a Get operation immediately after you create or modify a resource with one of the Put operations, the changes might not be returned. After a Get operation returns the most recent update, it always returns that updated resource until the resource is modified again. You can determine if an updated resource has been returned by looking at the checksum.

Runtime API Operations

Client applications use the following runtime API operations to communicate with Amazon Lex:

- [PostContent](#) – Takes speech or text input and returns intent information and a text or speech message to convey to the user. Currently, Amazon Lex supports the following audio formats:

Input audio formats – LPCM and Opus

Output audio formats – MPEG, OGG, and PCM

The `PostContent` operation supports audio input at 8 kHz and 16 kHz. Applications where the end user speaks with Amazon Lex over the telephone, such as an automated call center, can pass 8 kHz audio directly.

- [PostText](#) – Takes text as input and returns intent information and a text message to convey to the user.

Your client application uses the runtime API to call a specific Amazon Lex bot to process utterances — user text or voice input. For example, suppose that a user says "I want pizza." The client sends this user input to a bot using one of the Amazon Lex runtime API operations. From the user input, Amazon Lex recognizes that the user request is for the `OrderPizza` intent defined in the bot. Amazon Lex engages the user in a conversation to gather the required information, or slot data, such as pizza size, toppings, and number of pizzas. After the user provides all of the necessary slot data, Amazon Lex either invokes the Lambda function code hook to fulfill the intent, or returns the intent data to the client, depending on how the intent is configured.

Use the [PostContent](#) operation when your bot uses speech input. For example, an automated call center application can send speech to an Amazon Lex bot instead of an agent to address customer inquiries. You can use the 8 kHz audio format to send audio directly from the telephone to Amazon Lex.

The test window in the Amazon Lex console uses the [PostContent](#) API to send text and speech requests to Amazon Lex. You use this test window in the [Getting Started with Amazon Lex](#) exercises.

Lambda Functions As Code Hooks

You can configure your Amazon Lex bot to invoke a Lambda function as a code hook. The code hook can serve multiple purposes:

- Customizes the user interaction—For example, when Joe asks for available pizza toppings, you can use prior knowledge of Joe's choices to display a subset of toppings.
- Validates the user's input—Suppose that Jen wants to pick up flowers after hours. You can validate the time that Jen input and send an appropriate response.
- Fulfills the user's intent—After Joe provides all of the information for his pizza order, Amazon Lex can invoke a Lambda function to place the order with a local pizzeria.

When you configure an intent, you specify Lambda functions as code hooks in the following places:

- Dialog code hook for initialization and validation—This Lambda function is invoked on each user input, assuming Amazon Lex understood the user intent.
- Fulfillment code hook—This Lambda function is invoked after the user provides all of the slot data required to fulfill the intent.

You choose the intent and set the code hooks in the Amazon Lex console, as shown in the following screen shot:

OrderFlowers Latest ▾

▼ Sample utterances ⓘ

e.g. I would like to book a flight. +

I would like to pick up flowers ✕

I would like to order some flowers ✕

Order flowers ✕

▼ Lambda initialization and validation ⓘ

☒ Initialization and validation code hook

Lambda Function Name ▾

▼ Slots ⓘ

Priority	Required	Name	Slot type	Prompt
		e.g. Location	e.g. A... ▾	e.g. What city? ⚙️ +
1. ▾	☒	FlowerType	Flowe... ▾	1 ▾
2. ^ ▾	☒	PickupDate	AMA... ▾	Built-in ▾
3. ^ ▾	☒	PickupTime	AMA... ▾	Built-in ▾

What type of flow ⚙️ ✕

What day do you ⚙️ ✕

At what time do y ⚙️ ✕

▼ Confirmation prompt ⓘ

☒ Confirmation prompt

Confirm

Okay, your {FlowerType} will be ready for pickup by {Pickup} ⚙️

Cancel (if the user says "no")

Okay, I will not place your order. ⚙️

▼ Fulfillment ⓘ

☒ AWS Lambda function ☐ Return parameters to client

Lambda Function Name ▾

You can also set the code hooks using the `dialogCodeHook` and `fulfillmentActivity` fields in the [PutIntent](#) operation.

One Lambda function can perform initialization, validation, and fulfillment. The event data that the Lambda function receives has a field that identifies the caller as either a dialog or fulfillment code hook. You can use this information to run the appropriate portion of your code.

You can use a Lambda function to build a bot that can navigate complex dialogs. You use the `dialogAction` field in the Lambda function response to direct Amazon Lex to take specific actions. For example, you can use the `ElicitSlot` dialog action to tell Amazon Lex to ask the user for a slot value that isn't required. If you have a clarification prompt defined, you can use the `ElicitIntent` dialog action to elicit a new intent when the user is finished with the previous one.

For more information, see [Using Lambda Functions](#).

Managing Messages

Topics

- [Types of Messages](#)
- [Contexts for Configuring Messages](#)
- [Supported Message Formats](#)
- [Message Groups](#)
- [Response Cards](#)

When you create a bot, you can configure clarifying or informational messages that you want it to send to the client. Consider the following examples:

- You could configure your bot with the following clarification prompt:

I don't understand. What would you like to do?

Amazon Lex sends this message to the client if it doesn't understand the user's intent.

- Suppose that you create a bot to support an intent called `OrderPizza`. For a pizza order, you want users to provide information such as pizza size, toppings, and crust type. You could configure the following prompts:

```
What size pizza do you want?  
What toppings do you want?  
Do you want thick or thin crust?
```

After Amazon Lex determines the user's intent to order pizza, it sends these messages to the client to get information from the user.

This section explains designing user interactions in your bot configuration.

Types of Messages

A message can be a prompt or a statement.

- A *prompt* is typically a question and expects a user response.
- A *statement* is informational. It doesn't expect a response.

A message can include references to slot, session attributes, and request attributes. At runtime, Amazon Lex substitutes these references with actual values.

To refer to slots values that have been set, use the following syntax:

```
{SlotName}
```

To refer to session attributes, use the following syntax:

```
[SessionAttributeName]
```

To refer to request attributes, use the following syntax:

```
((RequestAttributeName))
```

Messages can include both slot values, session attributes and request attributes.

For example, suppose that you configure the following message in your bot's OrderPizza intent:

```
"Hey [FirstName], your {PizzaTopping} pizza will arrive in [DeliveryTime] minutes."
```

This message refers to both slot (PizzaTopping) and session attributes (FirstName and DeliveryTime). At runtime, Amazon Lex replaces these placeholders with values and returns the following message to the client:

```
"Hey John, your cheese pizza will arrive in 30 minutes."
```

To include brackets ([]) or braces ({} in a message, use the backslash (\) escape character. For example, the following message includes the curly braces and square brackets:

```
\{Text\} \[Text\]
```

The text returned to the client application looks like this:

```
{Text} [Text]
```

For information about session attributes, see the runtime API operations [PostText](#) and [PostContent](#). For an example, see [Book Trip](#).

Lambda functions can also generate messages and return them to Amazon Lex to send to the user. If you add Lambda functions when you configure your intent, you can create messages dynamically. By providing the messages while configuring your bot, you can eliminate the need to construct a prompt in your Lambda function.

Contexts for Configuring Messages


When you are creating your bot, you can create messages in different contexts, such as clarification prompts in bot, prompts for slot values, and messages from intents. Amazon Lex chooses an appropriate message in each context to return to your user. You can provide a group of messages for each context. If you do, Amazon Lex randomly chooses one message from the group. You can also specify the format of the message or group the messages together. For more information, see [Supported Message Formats](#).

If you have a Lambda function associated with an intent, you can override any of the messages that you configured at build time. A Lambda function is not required to use any of these messages, however.

Bot Messages

You can configure your bot with clarification prompts and session end messages. At runtime, Amazon Lex uses the clarification prompt if it doesn't understand the user's intent. You can

configure the number of times that Amazon Lex requests clarification before sending the session end message. You configure bot-level messages in the **Error Handling** section of the Amazon Lex console, as in the following image:



With the API, you configure messages by setting the `clarificationPrompt` and `abortStatement` fields in the [PutBot](#) operation.

If you use a Lambda function with an intent, the Lambda function might return a response directing Amazon Lex to ask a user's intent. If the Lambda function doesn't provide such a message, Amazon Lex uses the clarification prompt.

Slot Prompts

You must specify at least one prompt message for each of the required slots in an intent. At runtime, Amazon Lex uses one of these messages to prompt the user to provide a value for the slot. For example, for a `cityName` slot, the following is a valid prompt:

```
Which city would you like to fly to?
```


You can set one or more prompts for each slot using the console. You can also create groups of prompts using the [PutIntent](#) operation. For more information, see [Message Groups](#).

Responses


In the console, use the **Responses** section to build dynamic, engaging conversations for your bot. You can create one or more message groups for a response. At runtime, Amazon Lex builds a response by selecting one message from each message group. For more information about message groups, see [Message Groups](#).

For example, your first message group could contain different greetings: "Hello," "Hi," and "Greetings." The second message group could contain different forms of introduction: "I am the reservation bot" and "This is the reservation bot." A third message group could communicate the bot's capabilities: "I can help with car rentals and hotel reservations," "You can make car rentals and hotel reservations," and "I can help you rent a car and book a hotel."

Lex uses a message from each of the message groups to dynamically build the responses in a conversation. For example, one interaction could be the following:




Another one could be the following:




In either case, the user could respond with a new intent, such as the `BookCar` or `BookHotel` intent.

You can set up the bot to ask a follow-up question in the response. For example, for the preceding interaction, you could create a fourth message group with the following questions: "Can I help with a car or a hotel?", "Would you like to make a reservation now?", and "Is there anything that I can do for you?". For messages that include "No" as a response, you can create a follow-up prompt. The following image provides an example:



To create a follow-up prompt, choose **Wait for user reply**. Then type the message or messages that you want to send when the user says "No." When you create a response to use as a follow-up prompt, you must also specify an appropriate statement when the answer to the statement is "No." See the following image for an example:



To add responses to an intent with the API, use the `PutIntent` operation. To specify a response, set the `conclusionStatement` field in the `PutIntent` request. To set a follow-up prompt, set the `followUpPrompt` field and include the statement to send when the user says "No." You can't set both the `conclusionStatement` field and the `followUpPrompt` field on the same intent.

Supported Message Formats

When you use the [PostText](#) operation, or when you use the [PostContent](#) operation with the `Accept` header set to `text/plain; charset=utf8`, Amazon Lex supports messages in the following formats:

- **PlainText**—The message contains plain UTF-8 text.
- **SSML**—The message contains text formatted for voice output.
- **CustomPayload**—The message contains a custom format that you have created for your client. You can define the payload to meet the needs of your application.
- **Composite**—The message is a collection of messages, one from each message group. For more information about message groups, see [Message Groups](#).

By default, Amazon Lex returns any one of the messages defined for a particular prompt. For example, if you define five messages to elicit a slot value, Amazon Lex chooses one of the messages randomly and returns it to the client.

If you want Amazon Lex to return a specific type of message to the client in a run-time request, set the `x-amzn-lex:accept-content-types` request parameter. The response is limited to the type or types requested. If there is more than one message of the specified type, Amazon Lex returns one at random. For more information about the `x-amz-lex:accept-content-types` header, see [Setting the Response Type](#).

Message Groups

A *message group* is a set of suitable responses to a particular prompt. Use message groups when you want your bot to dynamically build the responses in a conversation. When Amazon Lex returns a response to the client application, it randomly chooses one message from each group. You can create a maximum of five message groups for each response. Each group can contain a maximum of five messages. For examples of creating message groups in the console, see [Responses](#).

To create a message group, you can use the console or you can use the [PutBot](#), [PutIntent](#), or [PutSlotType](#) operations to assign a group number to a message. If you don't create a message

group, or if you create only one message group, Amazon Lex sends a single message in the Message field. Client applications get multiple messages in a response only when you have created more than one message group in the console, or when you create more than one message group when you create or update an intent with the [PutIntent](#) operation.

When Amazon Lex sends a message from a group, the response's Message field contains an escaped JSON object that contains the messages. The following example shows the contents of the Message field when it contains multiple messages.

 Note

The example is formatted for readability. A response doesn't contain carriage returns (CR).

```
{\"messages\":[
  {\"type\": \"PlainText\", \"group\": 0, \"value\": \"Plain text\"},
  {\"type\": \"SSML\", \"group\": 1, \"value\": \"SSML text\"},
  {\"type\": \"CustomPayload\", \"group\": 2, \"value\": \"Custom payload\"}
]}
```

You can set the format of the messages. The format can be one of the following:

- PlainText—The message is in plain UTF-8 text.
- SSML—The message is Speech Synthesis Markup Language (SSML).
- CustomPayload—The message is in a custom format that you specified.

To control the format of messages that the PostContent and PostText operations return in the Message field, set the `x-amz-lex:accept-content-types` request attribute. For example, if you set the header to the following, you receive only plain text and SSML messages in the response:

```
x-amz-lex:accept-content-types: PlainText, SSML
```

If you request a specific message format and a message group doesn't contain that a message with that format, you get a `NoUsableMessageException` exception. When you use a message group to group messages by type, don't use the `x-amz-lex:accept-content-types` header.

For more information about the `x-amz-lex:accept-content-types` header, see [Setting the Response Type](#).


Response Cards

Note

Response cards do not work with Amazon Connect chat. However, see [Add interactive messages to chat](#) for similar functionality.

A *response card* contains a set of appropriate responses to a prompt. Use response cards to simplify interactions for your users and increase your bot's accuracy by reducing typographical errors in text interactions. You can send a response card for each prompt that Amazon Lex sends to your client application. You can use response cards with Facebook Messenger, Slack, Twilio, and your own client applications.

For example, in a taxi application, you can configure an option in the response card for "Home" and set the value to the user's home address. When the user selects this option, Amazon Lex receives the entire address as the input text. See the following image:



TAXI	
Where To?	Where would you like to go?
Home	
Work	
Someplace else	

You can define a response card for the following prompts:

- Conclusion statement
- Confirmation prompt
- Follow-up prompt

- Rejection statement
- Slot type utterances

You can define only one response card for each prompt.

You configure response cards when you create an intent. You can define a static response card at build time using the console or the [PutIntent](#) operation. Or you can define a dynamic response card at runtime in a Lambda function. If you define both static and dynamic response cards, the dynamic response card takes precedence.

Amazon Lex sends response cards in the format that the client understands. It transforms response cards for Facebook Messenger, Slack, and Twilio. For other clients, Amazon Lex sends a JSON structure in the [PostText](#) response. For example, if the client is Facebook Messenger, Amazon Lex transforms the response card to a generic template. For more information about Facebook Messenger generic templates, see [Generic Template](#) on the Facebook website. For an example of the JSON structure, see [Generating Response Cards Dynamically](#).

You can use response cards only with the [PostText](#) operation. You can't use response cards with the [PostContent](#) operation.

Defining Static Response Cards

Define static response cards with the [PutBot](#) operation or the Amazon Lex console when you create an intent. A static response card is defined at the same time as the intent. Use a static response card when the responses are fixed. Suppose that you are creating a bot with an intent that has a slot for flavor. When defining the flavor slot, you specify prompts, as shown in the following console screenshot:

Priority	Required	Name	Slot type	Prompt
			e.g. AMA...	e.g. What city?
1.	☒	teaSize	teaSize	What size iced tea wc
2.	☒	teaFlavor	teaFlavor	Would you like a flavo

When defining prompts, you can optionally associate a response card and define details with the [PutBot](#) operation, or, in the Amazon Lex console, as shown in the following example:

teaFlavor Prompts

maximum number of retries

2

Corresponding utterances

e.g. I would like to go to {toCity}

+

Prompt response cards

0

Card image

Card title

What Flavor?

Card subtitle

What flavor tea would

Preview

Facebook

Button value

lemon

Button title

Lemon

Raspberry

Plain

e.g. Button title

e.g. Button title

Delete card

What Flavor?

What flavor tea would you like?

Lemon


Raspberry

Plain

Cancel

Save

Now suppose that you've integrated your bot with Facebook Messenger. The user can click the buttons to choose a flavor, as shown in the following illustration:



To customize the content of a response card, you can refer to session attributes. At runtime, Amazon Lex substitutes these references with appropriate values from the session attributes. For more information, see [Setting Session Attributes](#). For an example, see [Using a Response Card](#).

Generating Response Cards Dynamically

To generate response cards dynamically at runtime, use the initialization and validation Lambda function for the intent. Use a dynamic response card when the responses are determined at runtime in the Lambda function. In response to user input, the Lambda function generates a response card and returns it in the `dialogAction` section of the response. For more information, see [Response Format](#).

The following is a partial response from a Lambda function that shows the `responseCard` element. It generates a user experience similar to the one shown in the preceding section.

```
responseCard: {  
  "version": 1,
```

```
"contentType": "application/vnd.amazonaws.card.generic",
"genericAttachments": [
  {
    "title": "What Flavor?",
    "subtitle": "What flavor do you want?",
    "imageUrl": "Link to image",
    "attachmentLinkUrl": "Link to attachment",
    "buttons": [
      {
        "text": "Lemon",
        "value": "lemon"
      },
      {
        "text": "Raspberry",
        "value": "raspberry"
      },
      {
        "text": "Plain",
        "value": "plain"
      }
    ]
  }
]
```

For an example, see [Schedule Appointment](#).

Managing Conversation Context

Conversation context is the information that a user, your application, or a Lambda function provides to an Amazon Lex bot to fulfill an intent. Conversation context includes slot data that the user provides, request attributes set by the client application, and session attributes that the client application and Lambda functions create.

Topics

- [Setting Intent Context](#)
- [Using Default Slot Values](#)
- [Setting Session Attributes](#)
- [Setting Request Attributes](#)
- [Setting the Session Timeout](#)

- [Sharing Information Between Intents](#)
- [Setting Complex Attributes](#)

Setting Intent Context

You can have Amazon Lex trigger intents based on *context*. A *context* is a state variable that can be associated with an intent when you define a bot.

You configure the contexts for an intent when you create the intent using the console or using the [PutIntent](#) operation. You can only use contexts in the English (US) (en-US) locale, and only if you set the `enableModelImprovements` parameter to `true` when you created the bot with the [PutBot](#) operation.

There are two types of relationships for contexts, output contexts and input contexts. An *output context* becomes active when an associated intent is fulfilled. An output context is returned to your application in the response from the [PostText](#) or [PostContent](#) operation, and it is set for the current session. After a context is activated, it stays active for the number of turns or time limit configured when the context was defined.

An *input context* specifies conditions under which an intent can be recognized. An intent can only be recognized during a conversation when all of its input contexts are active. An intent with no input contexts is always eligible for recognition.

Amazon Lex automatically manages the lifecycle of contexts that are activated by fulfilling intents with output contexts. You can also set active contexts in a call to the `PostContent` or `PostText` operation.

You can also set the context of a conversation using the Lambda function for the intent. The output context from Amazon Lex is sent to the Lambda function input event. The Lambda function can send contexts in its response. For more information, see [Lambda Function Input Event and Response Format](#).

For example, suppose you have an intent to book a rental car that is configured to return an output context called "book_car_fulfilled". When the intent is fulfilled, Amazon Lex sets the output context variable "book_car_fulfilled". Since "book_car_fulfilled" is an active context, an intent with the "book_car_fulfilled" context set as an input context is now considered for recognition, as long as a user utterance is recognized as an attempt to elicit that intent. You can use this for intents that only make sense after booking a car, such as emailing a receipt or modifying a reservation.


Output Context

Amazon Lex makes an intent's output contexts active when the intent is fulfilled. You can use the output context to control the intents eligible to follow up the current intent.

Each context has a list of parameters that are maintained in the session. The parameters are the slot values for the fulfilled intent. You can use these parameters to pre-populate slot values for other intents. For more information, see [Using Default Slot Values](#).

You configure the output context when you create an intent with the console or with the [PutIntent](#) operation. You can configure an intent with more than one output context. When the intent is fulfilled, all of the output contexts are activated and returned in the [PostText](#) or [PostContent](#) response.

The following shows assigning an output context to an intent using the console.



The screenshot shows the Amazon Lex console interface for configuring an intent. It features two sections: 'Input tags' and 'Output tags'. The 'Input tags' section has a dropdown menu currently showing 'Input context'. The 'Output tags' section also has a dropdown menu showing 'Output context'. Below the 'Output tags' dropdown, there is a tag named 'order_complete' with a close button (X). Underneath this tag, a gear icon indicates a configuration of '5 turns 90 secs'.

When you define an output context you also define its *time to live*, the length of time or number of turns that the context is included in responses from Amazon Lex. A *turn* is one request from your application to Amazon Lex. Once the number of turns or the time has expired, the context is no longer active.

Your application can use the output context as needed. For example, your application can use the output context to:

- Change the behavior of the application based on the context. For example, a travel application could have a different action for the context "book_car_fulfilled" than "rental_hotel_fulfilled."
- Return the output context to Amazon Lex as the input context for the next utterance. If Amazon Lex recognizes the utterance as an attempt to elicit an intent, it uses the context to limit the intents that can be returned to ones with the specified context.

Input Context

You set an input context to limit the points in the conversation where the intent is recognized. Intents without an input context are always eligible to be recognized.

You set the input contexts that an intent responds to using the console or the `PutIntent` operation. An intent can have more than one input context. The following shows assigning an input context to an intent using the console.

▼ Context ⓘ

Input tags ⓘ

▼

✕

Output tags ⓘ

▼

For an intent with more than one input context, all contexts must be active to trigger the intent. You can set an input context when you call the [PostText](#), [PostContent](#), or [PutSession](#) operation.

You can configure the slots in an intent to take default values from the current active context. Default values are used when Amazon Lex recognizes a new intent but doesn't receive a slot value. You specify the context name and slot name in the form `#context-name.parameter-name` when you define the slot. For more information, see [Using Default Slot Values](#).

Using Default Slot Values

When you use a default value, you specify a source for a slot value to be filled for new intents when no slot is provided by the user's input. This source can be previous dialog, request or session attributes, or a fixed value that you set at build-time.

You can use the following as the source for your default values.

- Previous dialog (contexts) – `#context-name.parameter-name`
- Session attributes – `[attribute-name]`
- Request attributes – `<attribute-name>`
- Fixed value – Any value that doesn't match the previous

When you use the [PutIntent](#) operation to add slots to an intent, you can add a list of default values. Default values are used in the order that they are listed. For example, suppose you have an intent with a slot with the following definition:

```
"slots": [  
  {  
    "name": "reservation-start-date",  
    "defaultValueSpec": {  
      "defaultValueList": [  
        {  
          "defaultValue": "#book-car-fulfilled.startDate"  
        },  
        {  
          "defaultValue": "[reservationStartDate]"  
        }  
      ]  
    },  
    Other slot configuration settings  
  }  
]
```

When the intent is recognized, the slot named "reservation-start-date" has its value set to one of the following.

1. If the "book-car-fulfilled" context is active, the value of the "startDate" parameter is used as the default value.
2. If the "book-car-fulfilled" context is not active, or if the "startDate" parameter is not set, the value of the "reservationStartDate" session attribute is used as the default value.
3. If neither of the first two default values are used, then the slot doesn't have a default value and Amazon Lex will elicit a value as usual.

If a default value is used for the slot, the slot is not elicited even if it is required.

Setting Session Attributes

Session attributes contain application-specific information that is passed between a bot and a client application during a session. Amazon Lex passes session attributes to all Lambda functions configured for a bot. If a Lambda function adds or updates session attributes, Amazon Lex passes the new information back to the client application. For example:

- In [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#), the example bot uses the price session attribute to maintain the price of flowers. The Lambda function sets this attribute based on the type of flowers that was ordered. For more information, see [Step 5 \(Optional\): Review the Details of the Information Flow \(Console\)](#).
- In [Book Trip](#), the example bot uses the currentReservation session attribute to maintain a copy of the slot type data during the conversation to book a hotel or to book a rental car. For more information, see [Details of the Information Flow](#).

Use session attributes in your Lambda functions to initialize a bot and to customize prompts and response cards. For example:

- Initialization — In a pizza ordering bot, the client application passes the user's location as a session attribute in the first call to the [PostContent](#) or [PostText](#) operation. For example, "Location": "111 Maple Street". The Lambda function uses this information to find the closest pizzeria to place the order.
- Personalize prompts — Configure prompts and response cards to refer to session attributes. For example, "Hey [FirstName], what toppings would you like?" If you pass the user's first name as a session attribute (`{"FirstName": "Jo"}`), Amazon Lex substitutes the name for the placeholder. It then sends a personalized prompt to the user, "Hey Jo, which toppings would you like?"

Session attributes persist for the duration of the session. Amazon Lex stores them in an encrypted data store until the session ends. The client can create session attributes in a request by calling either the [PostContent](#) or the [PostText](#) operation with the `sessionAttributes` field set to a value. A Lambda function can create a session attribute in a response. After the client or a Lambda function creates a session attribute, the stored attribute value is used any time that the client application doesn't include `sessionAttribute` field in a request to Amazon Lex.

For example, suppose you have two session attributes, `{"x": "1", "y": "2"}`. If the client calls the `PostContent` or `PostText` operation without specifying the `sessionAttributes` field, Amazon Lex calls the Lambda function with the stored session attributes (`{"x": 1, "y": 2}`). If the Lambda function doesn't return session attributes, Amazon Lex returns the stored session attributes to the client application.

If either the client application or a Lambda function passes session attributes, Amazon Lex updates the stored session attributes. Passing an existing value, such as `{"x": 2}`, updates the stored value. If you pass a new set of session attributes, such as `{"z": 3}`, the existing values are

removed and only the new value is kept. When an empty map, {}, is passed, stored values are erased.

To send session attributes to Amazon Lex, you create a string-to-string map of the attributes. The following shows how to map session attributes:

```
{
  "attributeName": "attributeValue",
  "attributeName": "attributeValue"
}
```

For the `PostText` operation, you insert the map into the body of the request using the `sessionAttributes` field, as follows:

```
"sessionAttributes": {
  "attributeName": "attributeValue",
  "attributeName": "attributeValue"
}
```

For the `PostContent` operation, you base64 encode the map, and then send it as the `x-amz-lex-session-attributes` header.

If you are sending binary or structured data in a session attribute, you must first transform the data to a simple string. For more information, see [Setting Complex Attributes](#).

Setting Request Attributes

Request attributes contain request-specific information and apply only to the current request. A client application sends this information to Amazon Lex. Use request attributes to pass information that doesn't need to persist for the entire session. You can create your own request attributes or you can use predefined attributes. To send request attributes, use the `x-amz-lex-request-attributes` header in a [the section called "PostContent"](#) or the `requestAttributes` field in a [the section called "PostText"](#) request. Because request attributes don't persist across requests like session attributes do, they are not returned in `PostContent` or `PostText` responses.

Note

To send information that persists across requests, use session attributes.

The namespace `x-amz-lex:` is reserved for the predefined request attributes. Don't create request attributes with the prefix `x-amz-lex:`.

Setting Predefined Request Attributes

Amazon Lex provides predefined request attributes for managing the way that it processes information sent to your bot. The attributes do not persist for the entire session, so you must send the predefined attributes in each request. All predefined attributes are in the `x-amz-lex:` namespace.

In addition to the following predefined attributes, Amazon Lex provides predefined attributes for messaging platforms. For a list of those attributes, see [Deploying an Amazon Lex Bot on a Messaging Platform](#).

Setting the Response Type

If you have two client applications that have different capabilities, you may need to limit the format of messages in a response. For example, you might want to restrict messages sent to a Web client to plain text, but enable a mobile client to use both plain text and Speech Synthesis Markup Language (SSML). To set the format of messages returned by the [PostContent](#) and [PostText](#) operations, use the `x-amz-lex:accept-content-types` request attribute.

You can set the attribute to any combination of the following message types:

- `PlainText`—The message contains plain UTF-8 text.
- `SSML`—The message contains text formatted for voice output.
- `CustomPayload`—The message contains a custom format that you have created for your client. You can define the payload to meet the needs of your application.

Amazon Lex returns only messages with the specified type in the `Message` field of the response. You can set more than one value by separating values with a comma. If you are using message groups, every message group must contain at least one message of the specified type. Otherwise, you get a `NoUsableMessageException` error. For more information, see [Message Groups](#).

Note

The `x-amz-lex:accept-content-types` request attribute has no effect on the contents of the HTML body. The contents of a `PostText` operation response is always plain UTF-8

text. The body of a `PostContent` operation response contains data in the format set in the `Accept` header in the request.

Setting the Preferred Time Zone

To set the time zone used to resolve dates so that it is relative to the user's time zone, use the `x-amz-lex:time-zone` request attribute. If you do not specify a time zone in the `x-amz-lex:time-zone` attribute, the default depends on the region that you are using for your bot.

Region	Default time zone
US East (N. Virginia)	America/New_York
US West (Oregon)	America/Los_Angeles
Asia Pacific (Singapore)	Asia/Singapore
Asia Pacific (Sydney)	Australia/Sydney
Asia Pacific (Tokyo)	Asia/Tokyo
Europe (Frankfurt)	Europe/Berlin
Europe (Ireland)	Europe/Dublin
Europe (London)	Europe/London

For example, if the user responds tomorrow in response to the prompt "Which day would you like your package delivered?" the actual *date* that the package is delivered depends on the user's time zone. For example, when it is 01:00 September 16 in New York, it is 22:00 September 15 in Los Angeles. If your service is running in the US East (N. Virginia) Region and a person in Los Angeles orders a package to be delivered "tomorrow" using the default time zone, the package would be delivered on the 17th, not the 16th. However, if you set the `x-amz-lex:time-zone` request attribute to `America/Los_Angeles`, the package would be delivered on the 16th.

You can set the attribute to any of the Internet Assigned Number Authority (IANA) time zone names. For the list of time zone names, see the [List of tz database time zones](#) on Wikipedia.

Setting User-Defined Request Attributes

A *user-defined request attribute* is data that you send to your bot in each request. You send the information in the `amz-lex-request-attributes` header of a `PostContent` request or in the `requestAttributes` field of a `PostText` request.

To send request attributes to Amazon Lex, you create a string-to-string map of the attributes. The following shows how to map request attributes:

```
{
  "attributeName": "attributeValue",
  "attributeName": "attributeValue"
}
```

For the `PostText` operation, you insert the map into the body of the request using the `requestAttributes` field, as follows:

```
"requestAttributes": {
  "attributeName": "attributeValue",
  "attributeName": "attributeValue"
}
```

For the `PostContent` operation, you base64 encode the map, and then send it as the `x-amz-lex-request-attributes` header.

If you are sending binary or structured data in a request attribute, you must first transform the data to a simple string. For more information, see [Setting Complex Attributes](#).

Setting the Session Timeout

Amazon Lex retains context information—slot data and session attributes—until a conversation session ends. To control how long a session lasts for a bot, set the session timeout. By default, session duration is 5 minutes, but you can specify any duration between 0 and 1,440 minutes (24 hours).

For example, suppose that you create a `ShoeOrdering` bot that supports intents such as `OrderShoes` and `GetOrderStatus`. When Amazon Lex detects that the user's intent is to order shoes, it asks for slot data. For example, it asks for shoe size, color, brand, etc. If the user provides some of the slot data but doesn't complete the shoe purchase, Amazon Lex remembers all of the

slot data and session attributes for the entire session. If the user returns to the session before it expires, he or she can provide the remaining slot data, and complete the purchase.

In the Amazon Lex console, you set the session timeout when you create a bot. With the AWS command line interface (AWS CLI) or API, you set the timeout when you create or update a bot with the [PutBot](#) operation by setting the [idleSessionTTLInSeconds](#) field.

Sharing Information Between Intents

Amazon Lex supports sharing information between intents. To share between intents, use session attributes.

For example, a user of the `ShoeOrdering` bot starts by ordering shoes. The bot engages in a conversation with the user, gathering slot data, such as shoe size, color, and brand. When the user places an order, the Lambda function that fulfills the order sets the `orderNumber` session attribute, which contains the order number. To get the status of the order, the user uses the `GetOrderStatus` intent. The bot can ask the user for slot data, such as order number and order date. When the bot has the required information, it returns the status of the order.

If you think that your users might switch intents during the same session, you can design your bot to return the status of the latest order. Instead of asking the user for order information again, you use the `orderNumber` session attribute to share information across intents and fulfill the `GetOrderStatus` intent. The bot does this by returning the status of the last order that the user placed.

For an example of cross-intent information sharing, see [Book Trip](#).

Setting Complex Attributes

Session and request attributes are string-to-string maps of attributes and values. In many cases, you can use the string map to transfer attribute values between your client application and a bot. In some cases, however, you might need to transfer binary data or a complex structure that can't be easily converted to a string map. For example, the following JSON object represents an array of the three most populous cities in the United States:

```
{
  "cities": [
    {
      "city": {
        "name": "New York",
```

```
        "state": "New York",
        "pop": "8537673"
    },
    {
        "city": {
            "name": "Los Angeles",
            "state": "California",
            "pop": "3976322"
        }
    },
    {
        "city": {
            "name": "Chicago",
            "state": "Illinois",
            "pop": "2704958"
        }
    }
]
```

This array of data doesn't translate well to a string-to-string map. In such a case, you can transform an object to a simple string so that you can send it to your bot with the [PostContent](#) and [PostText](#) operations.

For example, if you are using JavaScript, you can use the `JSON.stringify` operation to convert an object to JSON, and the `JSON.parse` operation to convert JSON text to a JavaScript object:

```
// To convert an object to a string.
var jsonString = JSON.stringify(object, null, 2);
// To convert a string to an object.
var obj = JSON.parse(JSON string);
```

To send session attributes with the `PostContent` operation, you must base64 encode the attributes before you add them to the request header, as shown in the following JavaScript code:

```
var encodedAttributes = new Buffer(attributeString).toString("base64");
```

You can send binary data to the `PostContent` and `PostText` operations by first converting the data to a base64-encoded string, and then sending the string as the value in the session attributes:

```
"sessionAttributes" : {  
  "binaryData": "base64 encoded data"  
}
```

Using Confidence Scores

When a user makes an utterance, Amazon Lex uses natural language understanding (NLU) to understand the user's request and return the proper intent. By default, Amazon Lex returns the most likely intent defined by your bot.

In some cases it may be difficult for Amazon Lex to determine the most likely intent. For example, the user might make an ambiguous utterance, or there might be two intents that are similar. To help determine the proper intent, you can combine your domain knowledge with the *confidence scores* of a list of alternative intents. A confidence score is a rating that Amazon Lex provides that shows how confident it is that an intent is the correct intent.

To determine the difference between two alternative intents, you can compare their confidence scores. For example, if one intent has a confidence score of 0.95 and another has a score of 0.65, the first intent is probably correct. However, if one intent has a score of 0.75 and another has a score of 0.72, there is ambiguity between the two intents that you may be able to discriminate using domain knowledge in your application.

You can also use confidence scores to create test applications that determine if changes to an intent's utterances make a difference in the behavior of the bot. For example, you can get the confidence scores for a bot's intents using a set of utterances, then update the intents with new utterances. You can then check the confidence scores to see if there was an improvement.

The confidence scores that Amazon Lex returns are comparative values. You should not rely on them as an absolute score. The values may change based on improvements to Amazon Lex.

When you use confidence scores, Amazon Lex returns the most likely intent and up to 4 alternative intents with their associated scores in each response. If all of the confidence scores are less than a threshold, Amazon Lex includes the `AMAZON.FallbackIntent`, the `AMAZON.KendraSearchIntent`, or both, if you have them configured. You can use the default threshold or you can set your own threshold.

The following JSON code shows the `alternativeIntents` field in the response from the [PostText](#) operation.

```
"alternativeIntents": [  
  {  
    "intentName": "string",  
    "nluIntentConfidence": {  
      "score": number  
    },  
    "slots": {  
      "string" : "string"  
    }  
  }  
],
```

Set the threshold when you create or update a bot. You can use either the API or the Amazon Lex console. For the regions listed below you need to opt-in to enable accuracy improvements and confidence scores. In the console, choose confidence scores in the **Advanced Options** section. Using the API, set the `enableModelImprovements` parameter when you call the [PutBot](#) operation. :

- US East (N. Virginia) (us-east-1)
- US West (Oregon) (us-west-2)
- Asia Pacific (Sydney) (ap-southeast-2)
- Europe (Ireland) (eu-west-1)

In all other regions, accuracy improvements and confidence score support is available by default.

To change the confidence threshold, set it in the console or using the [PutBot](#) operation. The threshold must be a number between 1.00 and 0.00.

To use the console, set the confidence threshold when you create or update your bot.

To set the confidence threshold when creating a bot (Console)

- On **Create your bot**, enter a value in the **Confidence score threshold** field.

To update the confidence threshold (Console)

1. From the list of your bots, choose the bot to update.
2. Choose the **Settings** tab.

3. In the left navigation, choose **General**.
4. Update the value in the **Confidence score threshold** field.

To set or update the confidence threshold (SDK)

- Set the `nluIntentConfidenceThreshold` parameter of the [PutBot](#) operation. The following JSON code shows the parameter being set.

```
"nluIntentConfidenceThreshold": 0.75,
```

Session Management

To change the intent that Amazon Lex uses in a conversation with the user, you can use the response from your dialog code hook Lambda function, or you can use the session management APIs in your custom application.

Using a Lambda function

When you use a Lambda function, Amazon Lex calls it with a JSON structure that contains the input to the function. The JSON structure contains a field called `currentIntent` that contains the intent that Amazon Lex has identified as the most likely intent for the user's utterance. The JSON structure also includes an `alternativeIntents` field that contains up to four additional intents that may satisfy the user's intent. Each intent includes a field called `nluIntentConfidenceScore` that contains the confidence score that Amazon Lex assigned to the intent.

To use an alternative intent, you specify it in the `ConfirmIntent` or the `ElicitSlot` dialog action in your Lambda function.

For more information, see [Using Lambda Functions](#).

Using the Session Management API

To use a different intent from the current intent, use the [PutSession](#) operation. For example, if you decide that the first alternative is preferable to the intent that Amazon Lex chose, you can use the `PutSession` operation to change intents so that the next intent that the user interacts with is the one that you selected.

For more information, see [Managing Sessions With the Amazon Lex API](#).

Conversation Logs

You enable *conversation logs* to store bot interactions. You can use these logs to review the performance of your bot and to troubleshoot issues with conversations. You can log text for the [PostText](#) operation. You can log both text and audio for the [PostContent](#) operation. By enabling conversation logs you get a detailed view of conversations that users have with your bot.

For example, a session with your bot has a session ID. You can use this ID to get the transcript of the conversation including user utterances and the corresponding bot responses. You also get metadata such as intent name and slot values for an utterance.

Note

You can't use conversation logs with a bot subject to the Children's Online Privacy Protection Act (COPPA).

Conversation logs are configured for an alias. Each alias can have different settings for their text and audio logs. You can enable text logs, audio logs, or both for each alias. Text logs store text input, transcripts of audio input, and associated metadata in CloudWatch Logs. Audio logs store audio input in Amazon S3. You can enable encryption of text and audio logs using AWS KMS customer managed CMKs.

To configure logging, use the console or the [PutBotAlias](#) operation. You can't log conversations for the \$LATEST alias of your bot or for the test bot available in the Amazon Lex console. After enabling conversation logs for an alias, [PostContent](#) or [PostText](#) operation for that alias logs the text or audio utterances in the configured CloudWatch Logs log group or S3 bucket.

Topics

- [IAM Policies for Conversation Logs](#)
- [Configuring Conversation Logs](#)
- [Encrypting Conversation Logs](#)
- [Viewing Text Logs in Amazon CloudWatch Logs](#)
- [Accessing Audio Logs in Amazon S3](#)
- [Monitoring Conversation Log Status with CloudWatch Metrics](#)

IAM Policies for Conversation Logs

Depending on the type of logging that you select, Amazon Lex requires permission to use Amazon CloudWatch Logs and Amazon Simple Storage Service (S3) buckets to store your logs. You must create AWS Identity and Access Management roles and permissions to enable Amazon Lex to access these resources.

Creating an IAM Role and Policies for Conversation Logs

To enable conversation logs, you must grant write permission for CloudWatch Logs and Amazon S3. If you enable object encryption for your S3 objects, you need to grant access permission to the AWS KMS keys used to encrypt the objects.

You can use the IAM AWS Management Console, the IAM API, or the AWS Command Line Interface to create the role and policies. These instructions use the AWS CLI to create the role and policies. For information about creating policies with the console, see [Creating policies on the JSON tab](#) in the *AWS Identity and Access Management User Guide*.

Note

The following code is formatted for Linux and MacOS. For Windows, replace the Linux line continuation character (`\`) with a caret (`^`).

To create an IAM role for conversation logs

1. Create a document in the current directory called **LexConversationLogsAssumeRolePolicyDocument.json**, add the following code to it, and save it. This policy document adds Amazon Lex as a trusted entity to the role. This allows Lex to assume the role to deliver logs to the resources configured for conversation logs.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "lex.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

```
    }  
  ]  
}
```

2. In the AWS CLI, run the following command to create the IAM role for conversation logs.

```
aws iam create-role \  
    --role-name role-name \  
    --assume-role-policy-document file://  
LexConversationLogsAssumeRolePolicyDocument.json
```

Next, create and attach a policy to the role that enables Amazon Lex to write to CloudWatch Logs.

To create an IAM policy for logging conversation text to CloudWatch Logs

1. Create a document in the current directory called **LexConversationLogsCloudWatchLogsPolicy.json**, add the following IAM policy to it, and save it.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "logs:CreateLogStream",  
        "logs:PutLogEvents"  
      ],  
      "Resource": "arn:aws:logs:region:account-id:log-group:log-group-name:"  
    }  
  ]  
}
```

2. In the AWS CLI, create the IAM policy that grants write permission to the CloudWatch Logs log group.

```
aws iam create-policy \  
    --policy-name cloudwatch-policy-name \  
    --policy-document file://LexConversationLogsCloudWatchLogsPolicy.json
```

3. Attach the policy to the IAM role that you created for conversation logs.

```
aws iam attach-role-policy \  
  --policy-arn arn:aws:iam::account-id:policy/cloudwatch-policy-name \  
  --role-name role-name
```

If you are logging audio to an S3 bucket, create a policy that enables Amazon Lex to write to the bucket.

To create an IAM policy for audio logging to an S3 bucket

1. Create a document in the current directory called **LexConversationLogsS3Policy.json**, add the following policy to it, and save it.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "s3:PutObject"  
      ],  
      "Resource": "arn:aws:s3:::bucket-name/*"  
    }  
  ]  
}
```

2. In the AWS CLI, create the IAM policy that grants write permission to your S3 bucket.

```
aws iam create-policy \  
  --policy-name s3-policy-name \  
  --policy-document file://LexConversationLogsS3Policy.json
```

3. Attach the policy to the role that you created for conversation logs.

```
aws iam attach-role-policy \  
  --policy-arn arn:aws:iam::account-id:policy/s3-policy-name \  
  --role-name role-name
```

Granting Permission to Pass an IAM Role

When you use the console, the AWS Command Line Interface, or an AWS SDK to specify an IAM role to use for conversation logs, the user specifying the conversation logs IAM role must have permission to pass the role to Amazon Lex. To allow the user to pass the role to Amazon Lex, you must grant `PassRole` permission to the user, role, or group.

The following policy defines the permission to grant to the user, role, or group. You can use the `iam:AssociatedResourceArn` and `iam:PassedToService` condition keys to limit the scope of the permission. For more information, see [Granting a User Permissions to Pass a Role to an AWS Service](#) and [IAM and AWS STS Condition Context Keys](#) in the *AWS Identity and Access Management User Guide*.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::account-id:role/role-name",
      "Condition": {
        "StringEquals": {
          "iam:PassedToService": "lex.amazonaws.com"
        },
        "StringLike": {
          "iam:AssociatedResourceARN": "arn:aws:lex:region:account-id:bot:bot-name:bot-alias"
        }
      }
    }
  ]
}
```

Configuring Conversation Logs

You enable and disable conversation logs using the console or the `conversationLogs` field of the `PutBotAlias` operation. You can turn on or turn off audio logs, text logs, or both. Logging starts on new bot sessions. Changes to log settings aren't reflected for active sessions.

To store text logs, use an Amazon CloudWatch Logs log group in your AWS account. You can use any valid log group. The log group must be in the same region as the Amazon Lex bot. For more

information about creating a CloudWatch Logs log group, see [Working with Log Groups and Log Streams](#) in the *Amazon CloudWatch Logs User Guide*.

To store audio logs, use an Amazon S3 bucket in your AWS account. You can use any valid S3 bucket. The bucket must be in the same region as the Amazon Lex bot. For more information about creating an S3 bucket, see [Create a Bucket](#) in the *Amazon Simple Storage Service Getting Started Guide*.

You must provide an IAM role with policies that enable Amazon Lex to write to the configured log group or bucket. For more information, see [Creating an IAM Role and Policies for Conversation Logs](#).

If you create a service-linked role using the AWS Command Line Interface, you must add a custom suffix to the role using the `custom-suffix` option as follows:

```
aws iam create-service-linked-role \
  --aws-service-name lex.amazon.aws.com \
  --custom-suffix suffix
```

The IAM role that you use to enable conversation logs must have the `iam:PassRole` permission. The following policy should be attached to the role.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::account:role/role"
    }
  ]
}
```

Enabling Conversation Logs

To turn on logs using the console

1. Open the Amazon Lex console <https://console.aws.amazon.com/lex>.
2. From the list, choose a bot.
3. Choose the **Settings** tab, and then from the left menu choose **Conversation logs**.

4. In the list of aliases, choose the settings icon for the alias for which you want to configure conversation logs.
5. Select whether to log text, audio, or both.
6. For text logging, enter the Amazon CloudWatch Logs log group name.
7. For audio logging, enter the S3 bucket information.
8. Optional. To encrypt audio logs, choose the AWS KMS key to use for encryption.
9. Choose an IAM role with the required permissions.
10. Choose **Save** to start logging conversations.

To turn on text logs using the API

1. Call the [PutBotAlias](#) operation with an entry in the `logSettings` member of the `conversationLogs` field
 - Set the `destination` member to `CLOUDWATCH_LOGS`
 - Set the `logType` member to `TEXT`
 - Set the `resourceArn` member to the Amazon Resource Name (ARN) of the CloudWatch Logs log group that is the destination for the logs
2. Set the `iamRoleArn` member of the `conversationLogs` field to the Amazon Resource Name (ARN) of an IAM role that has the required permissions for enabling conversation logs on the specified resources.

To turn on audio logs using the API

1. Call the [PutBotAlias](#) operation with an entry in the `logSettings` member of the `conversationLogs` field
 - Set the `destination` member to `S3`
 - Set the `logType` member to `AUDIO`
 - Set the `resourceArn` member to the ARN of the Amazon S3 bucket where the audio logs are stored
 - Optional. To encrypt audio logs with a specific AWS KMS key, set the `kmsKeyArn` member of the ARN of the key that is used for encryption.

2. Set the `iamRoleArn` member of the `conversationLogs` field to the Amazon Resource Name (ARN) of an IAM role that has the required permissions for enabling conversation logs on the specified resources.

Disabling Conversation Logs

To turn off logs using the console

1. Open the Amazon Lex console <https://console.aws.amazon.com/lex>.
2. From the list, choose a bot.
3. Choose the **Settings** tab, and then from the left menu choose **Conversation logs**.
4. In the list of aliases, choose the settings icon for the alias for which you want to configure conversation logs.
5. Clear the check from text, audio, or both to turn off logging.
6. Choose **Save** to stop logging conversations.

To turn off logs using the API

- Call the `PutBotAlias` operation without the `conversationLogs` field.

To turn off text logs using the API

- If you are logging audio
 - Call the [PutBotAlias](#) operation with a `logSettings` entry only for AUDIO.
 - The call to the `PutBotAlias` operation must not have a `logSettings` entry for TEXT.
- If you are not logging audio
 - Call the [PutBotAlias](#) operation without the `conversationLogs` field.

To turn off audio logs using the API

- If you are logging text
 - Call the [PutBotAlias](#) operation with a `logSettings` entry only for TEXT.
 - The call to the `PutBotAlias` operation must not have a `logSettings` entry for AUDIO.
- If you are not logging text

- Call the [PutBotAlias](#) operation without the `conversationLogs` field.

Encrypting Conversation Logs

You can use encryption to help protect the contents of your conversation logs. For text and audio logs, you can use AWS KMS customer managed CMKs to encrypt data in your CloudWatch Logs log group and S3 bucket.

Note

Amazon Lex supports only symmetric CMKs. Do not use an asymmetric CMK to encrypt your data.

You enable encryption using an AWS KMS key on the CloudWatch Logs log group that Amazon Lex uses for text logs. You can't provide an AWS KMS key in the log settings to enable AWS KMS encryption of your log group. For more information, see [Encrypt Log Data in CloudWatch Logs Using AWS KMS](#) in the *Amazon CloudWatch Logs User Guide*.

For audio logs you use default encryption on your S3 bucket or specify an AWS KMS key to encrypt your audio objects. Even if your S3 bucket uses default encryption you can still specify a different AWS KMS key to encrypt your audio objects. For more information, see [Amazon S3 Default Encryption for S3 Buckets](#) in the *Amazon Simple Storage Service Developer Guide*.

Amazon Lex requires AWS KMS permissions if you choose to encrypt your audio logs. You need to attach additional policies to the IAM role used for conversation logs. If you use default encryption on your S3 bucket, your policy must grant access to the AWS KMS key configured for that bucket. If you specify an AWS KMS key in your audio log settings, you must grant access to that key.

If you have not created a role for conversation logs, see [IAM Policies for Conversation Logs](#).

To create an IAM policy for using an AWS KMS key for encrypting audio logs

1. Create a document in the current directory called `LexConversationLogsKMSPolicy.json`, add the following policy to it, and save it.

```
{  
  "Version": "2012-10-17",
```

```
"Statement": [  
  {  
    "Effect": "Allow",  
    "Action": [  
      "kms:GenerateDataKey"  
    ],  
    "Resource": "kms-key-arn"  
  }  
]
```

2. In the AWS CLI, create the IAM policy that grants permission to use the AWS KMS key for encrypting audio logs.

```
aws iam create-policy \  
  --policy-name kms-policy-name \  
  --policy-document file://LexConversationLogsKMSPolicy.json
```

3. Attach the policy to the role that you created for conversation logs.

```
aws iam attach-role-policy \  
  --policy-arn arn:aws:iam::account-id:policy/kms-policy-name \  
  --role-name role-name
```

Viewing Text Logs in Amazon CloudWatch Logs

Amazon Lex stores text logs for your conversations in Amazon CloudWatch Logs. To view the logs, you can use the CloudWatch Logs console or API. For more information, see [Search Log Data Using Filter Patterns](#) and [CloudWatch Logs Insights Query Syntax](#) in the *Amazon CloudWatch Logs User Guide*.

To view logs using the Amazon Lex console

1. Open the Amazon Lex console <https://console.aws.amazon.com/lex>.
2. From the list, choose a bot.
3. Choose the **Settings** tab, then from the left menu choose **Conversation logs**.
4. Choose the link under **Text logs** to view the logs for the alias in the CloudWatch console.

You can also use the CloudWatch console or API to view your log entries. To find the log entries, navigate to the log group that you configured for the alias. You find the log stream prefix for your logs in the Amazon Lex console or by using the [GetBotAlias](#) operation.

Log entries for a user utterance is in multiple log streams. An utterance in the conversation has an entry in one of the log streams with the specified prefix. An entry in the log stream contains the following information.

```
{
  "messageVersion": "1.0",
  "botName": "bot name",
  "botAlias": "bot alias",
  "botVersion": "bot version",
  "inputTranscript": "text used to process the request",
  "botResponse": "response from the bot",
  "intent": "matched intent",
  "nluIntentConfidence": "number",
  "slots": {
    "slot name": "slot value",
    "slot name": null,
    "slot name": "slot value"
    ...
  },
  "alternativeIntents": [
    {
      "name": "intent name",
      "nluIntentConfidence": "number",
      "slots": {
        "slot name": slot value,
        "slot name": null,
        "slot name": slot value
        ...
      }
    },
    {
      "name": "intent name",
      "nluIntentConfidence": number,
      "slots": {}
    }
  ],
  "developerOverride": "true" | "false",
  "missedUtterance": true | false,
  "inputDialogMode": "Text" | "Speech",
}
```

```

"requestId": "request ID",
"s3PathForAudio": "S3 path to audio file",
"userId": "user ID",
"sessionId": "session ID",
"sentimentResponse": {
    "sentimentScore": "{Positive: number, Negative: number, Neutral: number,
Mixed: number}",
    "sentimentLabel": "Positive" | "Negative" | "Neutral" | "Mixed"
},
"slotToElicit": "slot name",
"dialogState": "ElicitIntent" | "ConfirmIntent" | "ElicitSlot" | "Fulfilled" |
"ReadyForFulfillment" | "Failed",
"responseCard": {
    "genericAttachments": [
        ...
    ],
    "contentType": "application/vnd.amazonaws.card.generic",
    "version": 1
},
"locale": "locale",
"timestamp": "ISO 8601 UTC timestamp",
"kendraResponse": {
    "totalNumberOfResults": number,
    "resultItems": [
        {
            "id": "query ID",
            "type": "DOCUMENT" | "QUESTION_ANSWER" | "ANSWER",
            "additionalAttributes": [
                {
                    ...
                }
            ],
            "documentId": "document ID",
            "documentTitle": {
                "text": "title",
                "highlights": null
            },
            "documentExcerpt": {
                "text": "text",
                "highlights": [
                    {
                        "beginOffset": number,
                        "endOffset": number,
                        "topAnswer": true | false
                    }
                ]
            }
        }
    ]
}

```

```

        }
      ]
    },
    "documentURI": "URI",
    "documentAttributes": []
  }
],
"facetResults": [],
"sdkResponseMetadata": {
  "requestId": "request ID"
},
"sdkHttpMetadata": {
  "httpHeaders": {
    "Content-Length": "number",
    "Content-Type": "application/x-amz-json-1.1",
    "Date": "date and time",
    "x-amzn-RequestId": "request ID"
  },
  "httpStatusCode": 200
},
"queryId": "query ID"
},
"sessionAttributes": {
  "attribute name": "attribute value"
  ...
},
"requestAttributes": {
  "attribute name": "attribute value"
  ...
}
}

```

The contents of the log entry depends on the result of a transaction and the configuration of the bot and request.

- The `intent`, `slots`, and `slotToElicit` fields don't appear in an entry if the `missedUtterance` field is `true`.
- The `s3PathForAudio` field doesn't appear if audio logs are disabled or if the `inputDialogMode` field is `Text`.
- The `responseCard` field only appears when you have defined a response card for the bot.

- The `requestAttributes` map only appears if you have specified request attributes in the request.
- The `kendraResponse` field is only present when the `AMAZON.KendraSearchIntent` makes a request to search an Amazon Kendra index.
- The `developerOverride` field is true when an alternative intent was specified in the bot's Lambda function.
- The `sessionAttributes` map only appears if you have specified session attributes in the request.
- The `sentimentResponse` map only appears if you configure the bot to return sentiment values.

Note

The input format may change without a corresponding change in the `messageVersion`. Your code should not throw an error if new fields are present.

You must have a role and policy set to enable Amazon Lex to write to CloudWatch Logs. For more information see [IAM Policies for Conversation Logs](#).

Accessing Audio Logs in Amazon S3

Amazon Lex stores audio logs for your conversations in an S3 bucket.

To access audio logs using the console

1. Open the Amazon Lex console <https://console.aws.amazon.com/lex>.
2. From the list, choose a bot.
3. Choose the **Settings** tab, then from the left menu choose **Conversation logs**.
4. Choose the link under **Audio logs** to access the logs for the alias in the Amazon S3 console.

You can also use the Amazon S3 console or API to access audio logs. You can see the S3 object key prefix of the audio files in the Amazon Lex console, or in the `resourcePrefix` field in the `GetBotAlias` operation response.

Monitoring Conversation Log Status with CloudWatch Metrics

Use Amazon CloudWatch to monitor delivery metrics of your conversation logs. You can set alarms on metrics so that you are aware of issues with logging if they should occur.

Amazon Lex provides four metrics in the AWS/Lex namespace for conversation logs:

- `ConversationLogsAudioDeliverySuccess`
- `ConversationLogsAudioDeliveryFailure`
- `ConversationLogsTextDeliverySuccess`
- `ConversationLogsTextDeliveryFailure`

For more information, see [CloudWatch Metrics for Conversation Logs](#).

The success metrics show that Amazon Lex has successfully written your audio or text logs to their destinations.

The failure metrics show that Amazon Lex couldn't deliver audio or text logs to the specified destination. Typically, this is a configuration error. When your failure metrics are above zero, check the following:

- Make sure that Amazon Lex is a trusted entity for the IAM role.
- For text logging, make sure that the CloudWatch Logs log group exists. For audio logging, make sure that the S3 bucket exists.
- Make sure that the IAM role that Amazon Lex uses to access the CloudWatch Logs log group or S3 bucket has write permission for the log group or bucket.
- Make sure that the S3 bucket exists in the same region as the Amazon Lex bot and belongs to your account.
- If you are using an AWS KMS key for S3 encryption, make sure that there are no policies that prevent Amazon Lex from using your key and make sure that the IAM role you provide has the necessary AWS KMS permissions. For more information, see [IAM Policies for Conversation Logs](#).

Managing Sessions With the Amazon Lex API

When a user starts a conversation with your bot, Amazon Lex creates a *session*. The information exchanged between your application and Amazon Lex makes up the session state for the

conversation. When you make a request, the session is identified by a combination of the bot name and a user identifier that you specify. For more information about the user identifier, see the `userId` field in the [PostContent](#) or [PostText](#) operation.

The response from a session operation includes a unique session identifier that identifies a specific session with a user. You can use this identifier during testing or to help troubleshoot your bot.

You can modify the session state sent between your application and your bot. For example, you can create and modify session attributes that contain custom information about the session, and you can change the flow of the conversation by setting the dialog context to interpret the next utterance.

There are two ways that you can update session state. The first is to use a Lambda function with the `PostContent` or `PostText` operation that is called after each turn of the conversation. For more information, see [Using Lambda Functions](#). The other is to use the Amazon Lex runtime API in your application to make changes to the session state.

The Amazon Lex runtime API provides operations that enable you to manage session information for a conversation with your bot. The operations are the [PutSession](#) operation, the [GetSession](#) operation, and the [DeleteSession](#) operation. You use these operations to get information about the state of your user's session with your bot, and to have fine-grained control over the state.

Use the `GetSession` operation when you want to get the current state of the session. The operation returns the current state of the session, including the state of the dialog with your user, any session attributes that have been set and slot values for the last three intents that the user interacted with.

The `PutSession` operation enables you to directly manipulate the current session state. You can set the type of dialog action that the bot will perform next. This gives you control over the flow of the conversation with the bot. Set the `dialogActionType` field to `Delegate` to have Amazon Lex determine the next action for the bot.

You can use the `PutSession` operation to create a new session with a bot and set the intent that the bot should start with. You can also use the `PutSession` operation to change from one intent to another. When you create a session or change the intent you also can set session state, such as slot values and session attributes. When the new intent is finished, you have the option of restarting the prior intent. You can use the `GetSession` operation to get the dialog state of the prior intent from Amazon Lex and use the information to set the dialog state of the intent.

The response from the `PutSession` operation contains the same information as the `PostContent` operation. You can use this information to prompt the user for the next piece of information, just as you would with the response from the `PostContent` operation.

Use the `DeleteSession` operation to remove an existing session and start over with a new session. For example, when you are testing your bot you can use the `DeleteSession` operation to remove test sessions from your bot.

The session operations work with your fulfillment Lambda functions. For example, if your Lambda function returns `Failed` as the fulfillment state you can use the `PutSession` operation to set the dialog action type to `close` and fulfillmentState to `ReadyForFulfillment` to retry the fulfillment step.

Here are some things that you can do with the session operations:

- Have the bot start a conversation instead of waiting for the user.
- Switch intents during a conversation.
- Return to a previous intent.
- Start or restart a conversation in the middle of the interaction.
- Validate slot values and have the bot re-prompt for values that are not valid.

Each of these are described further below.

Switching Intents

You can use the `PutSession` operation to switch from one intent to another. You can also use it to switch back to a previous intent. You can use the `PutSession` operation to set session attributes or slot values for the new intent.

- Call the `PutSession` operation. Set the intent name to the name of the new intent and set the dialog action to `Delegate`. You can also set any slot values or session attributes required for the new intent.
- Amazon Lex will start a conversation with the user using the new intent.

Resuming a Prior Intent

To resume a prior intent you use the `GetSession` operation to get the summary of the intent, and then use the `PutSession` operation to set the intent to its previous dialog state.

- Call the `GetSession` operation. The response from the operation includes a summary of the dialog state of the last three intents that the user interacted with.
- Using the information from the intent summary, call the `PutSession` operation. This will return the user to the previous intent in the same place in the conversation.

In some cases it may be necessary to resume your user's conversation with your bot. For example, say that you have created a customer service bot. Your application determines that the user needs to talk to a customer service representative. After talking to the user, the representative can direct the conversation back to the bot with the information that they collected.

To resume a session, use steps similar to these:

- Your application determines that the user needs to speak to a customer service representative.
- Use the `GetSession` operation to get the current dialog state of the intent.
- The customer service representative talks to the user and resolves the issue.
- Use the `PutSession` operation to set the dialog state of the intent. This may include setting slot values, setting session attributes, or changing the intent.
- The bot resumes the conversation with the user.

You can use the `PutSession` operation `checkpointLabel` parameter to label an intent so that you can find it later. For example, a bot that asks a customer for information might go into a `Waiting` intent while the customer gathers the information. The bot creates a checkpoint label for the current intent and then starts the `Waiting` intent. When the customer returns the bot can find the previous intent using the checkpoint label and switch back.

The intent must be present in the `recentIntentSummaryView` structure returned by the `GetSession` operation. If you specify a checkpoint label in the `GetSession` operation request, it will return a maximum of three intents with that checkpoint label.

- Use the `GetSession` operation to get the current state of the session.

- Use the `PutSession` operation to add a checkpoint label to the last intent. If necessary you can use this `PutSession` call to switch to a different intent.
- When it is time to switch back to the labeled intent, call the `GetSession` operation to return a recent intent list. You can use the `checkpointLabelFilter` parameter so that Amazon Lex returns only intents with the specified checkpoint label.

Starting a New Session

If you want to have the bot start the conversation with your user, you can use the `PutSession` operation.

- Create a welcome intent with no slots and a conclusion message that prompts the user to state an intent. For example, "What would you like to order? You can say 'Order a drink' or 'Order a pizza.'"
- Call the `PutSession` operation. Set the intent name to the name of your welcome intent and set the dialog action to `Delegate`.
- Amazon Lex will respond with the prompt from your welcome intent to start the conversation with your user.

Validating Slot Values

You can validate responses to your bot using your client application. If the response isn't valid, you can use the `PutSession` operation to get a new response from your user. For example, suppose that your flower ordering bot can only sell tulips, roses, and lilies. If the user orders carnations, your application can do the following:

- Examine the slot value returned from the `PostText` or `PostContent` response.
- If the slot value is not valid, call the `PutSession` operation. Your application should clear the slot value, set the `slotToElicit` field, and set the `dialogAction.type` value to `elicitSlot`. Optionally, you can set the `message` and `messageFormat` fields if you want to change the message that Amazon Lex uses to elicit the slot value.

Bot Deployment Options

Currently, Amazon Lex provides the following bot deployment options:

- [AWS Mobile SDK](#) – You can build mobile applications that communicate with Amazon Lex using the AWS Mobile SDKs.
- Facebook Messenger – You can integrate your Facebook Messenger page with your Amazon Lex bot so that end users on Facebook can communicate with the bot. In the current implementation, this integration supports only text input messages.
- Slack – You can integrate your Amazon Lex bot with a Slack messaging application.
- Twilio – You can integrate your Amazon Lex bot with the Twilio Simple Messaging Service (SMS).

For examples, see [Deploying Amazon Lex Bots](#).

Built-in Intents and Slot Types

To make it easier to create bots, Amazon Lex allows you to use standard built-in intents and slot types.

Topics

- [Built-in Intents](#)
- [Built-in Slot Types](#)

Built-in Intents

For common actions, you can use the standard built-in intents library. To create an intent from a built-in intent, choose a built-intent in the console, and give it a new name. The new intent has the configuration of the base intent, such as the sample utterances.

In the current implementation, you can't do the following:

- Add or remove sample utterances from the base intent
- Configure slots for built-in intents

To add a built-in intent to a bot

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose the bot to add the built-in intent to.
3. In the navigation pane, choose the plus (+) next to **Intents**.

4. For **Add intent**, choose **Search existing intents**.
5. In the **Search intents** box, type the name of the built-in intent to add to your bot.
6. For **Copy built-in intent**, give the intent a name, and then choose **Add**.
7. Configure the intent as required for your bot.

Topics

- [AMAZON.CancelIntent](#)
- [AMAZON.FallbackIntent](#)
- [AMAZON.HelpIntent](#)
- [AMAZON.KendraSearchIntent](#)
- [AMAZON.PauseIntent](#)
- [AMAZON.RepeatIntent](#)
- [AMAZON.ResumeIntent](#)
- [AMAZON.StartOverIntent](#)
- [AMAZON.StopIntent](#)

Note

For the English (US) (en-US) locale, Amazon Lex supports intents from the Alexa standard built-in intents. For a list of built-in intents, see [Standard Built-in Intents](#) in the *Alexa Skills Kit*.

Amazon Lex doesn't support the following intents:

- AMAZON.YesIntent
- AMAZON.NoIntent
- The intents in the [Built-in Intent Library](#) in the *Alexa Skills Kit*

AMAZON.CancelIntent

Responds to words and phrases that indicate the user wants to cancel the current interaction. Your application can use this intent to remove slot type values and other attributes before ending the interaction with the user.

Common utterances:

- cancel
- never mind
- forget it

AMAZON.FallbackIntent

When a user's input to an intent isn't what a bot expects, you can configure Amazon Lex to invoke a *fallback intent*. For example, if the user input "I'd like to order candy" doesn't match an intent in your `OrderFlowers` bot, Amazon Lex invokes the fallback intent to handle the response.

You add a fallback intent by adding the built-in `AMAZON.FallbackIntent` intent type to your bot. You can specify the intent using the [PutBot](#) operation or by choosing the intent from the list of built-in intents in the console.

Invoking a fallback intent uses two steps. In the first step the fallback intent is matched based on the input from the user. When the fallback intent is matched, the way the bot behaves depends on the number of retries configured for a prompt. For example, if the maximum number of attempts to determine an intent is 2, the bot returns the bot's clarification prompt twice before invoking the fallback intent.

Amazon Lex matches the fallback intent in these situations:

- The user's input to an intent doesn't match the input that the bot expects
- Audio input is noise, or text input isn't recognized as words.
- The user's input is ambiguous and Amazon Lex can't determine which intent to invoke.

The fallback intent is invoked when:

- The bot doesn't recognize the user input as an intent after the configured number of tries for clarification when the conversation is started.
- An intent doesn't recognize the user input as a slot value after the configured number of tries.
- An intent doesn't recognize the user input as a response to a confirmation prompt after the configured number of tries.

You can use the following with a fallback intent:

- A fulfillment Lambda function
- A conclusion statement
- A follow up prompt

You can't add the following to a fallback intent:

- Utterances
- Slots
- An initialization and validation Lambda function
- A confirmation prompt

If you have configured both a cancel statement and a fallback intent for a bot, Amazon Lex uses the fallback intent. If you need your bot to have a cancel statement, you can use the fulfillment function for the fallback intent to provide the same behavior as a cancel statement. For more information, see the `abortStatement` parameter of the [PutBot](#) operation.

Using Clarification Prompts

If you provide your bot with a clarification prompt, the prompt is used to solicit a valid intent from the user. The clarification prompt will be repeated the number of times that you configured. After that the fallback intent will be invoked.

If you don't set a clarification prompt when you create a bot and the user doesn't start the conversation with a valid intent, Amazon Lex immediately calls your fallback intent .

When you use a fallback intent without a clarification prompt, Amazon Lex doesn't call the fallback under these circumstances:

- When the user responds to a follow-up prompt but doesn't provide an intent. For example, in response to a follow-up prompt that says "Would you like anything else today?", the user says "Yes." Amazon Lex returns a 400 Bad Request exception because it doesn't have a clarification prompt to send to the user to get an intent.
- When using an AWS Lambda function, you return an `ElicitIntent` dialog type. Because Amazon Lex doesn't have a clarification prompt to get an intent from the user, it returns a 400 Bad Request exception.

- When using the `PutSession` operation, you send an `ElicitIntent` dialog type. Because Amazon Lex doesn't have a clarification prompt to get an intent from the user, it returns a 400 Bad Request exception.

Using a Lambda Function with a Fallback Intent

When a fallback intent is invoked, the response depends on the setting of the `fulfillmentActivity` parameter to the [PutIntent](#) operation. The bot does one of the following:

- Returns the intent information to the client application.
- Calls the fulfillment Lambda function. It calls the function with the session variables that are set for the session.

For more information about setting the response when a fallback intent is invoked, see the `fulfillmentActivity` parameter of the [PutIntent](#) operation.

If you use the fulfillment Lambda function in your fallback intent, you can use this function to call another intent or to perform some form of communication with the user, such as collecting a callback number or opening a session with a customer service representative.

You can perform any action in a fallback intent Lambda function that you can perform in the fulfillment function for any other intent. For more information about creating a fulfillment function using AWS Lambda, see [Using Lambda Functions](#).

A fallback intent can be invoked multiple times in the same session. For example, suppose that your Lambda function uses the `ElicitIntent` dialog action to prompt the user for a different intent. If Amazon Lex can't infer the user's intent after the configured number of tries, it invokes the fallback intent again. It also invokes the fallback intent when the user doesn't respond with a valid slot value after the configured number of tries.

You can configure a Lambda function to keep track of the number of times that the fallback intent is called using a session variable. Your Lambda function can take a different action if it is called more times than the threshold that you set in your Lambda function. For more information about session variables, see [Setting Session Attributes](#).

AMAZON.HelpIntent

Responds to words or phrases that indicate the user needs help while interacting with your bot. When this intent is invoked, you can configure your Lambda function or application to provide

information about the your bot's capabilities, ask follow up questions about areas of help, or hand the interaction over to a human agent.

Common utterances:

- help
- help me
- can you help me

AMAZON.KendraSearchIntent

To search documents that you have indexed with Amazon Kendra, use the `AMAZON.KendraSearchIntent` intent. When Amazon Lex can't determine the next action in a conversation with the user, it triggers the search intent.

The `AMAZON.KendraSearchIntent` is available only in the English (US) (en-US) locale and in the US East (N. Virginia), US West (Oregon) and Europe (Ireland) Regions.

Amazon Kendra is a machine-learning-based search service that indexes natural language documents such as PDF documents or Microsoft Word files. It can search indexed documents and return the following types of responses to a question:

- An answer
- An entry from a FAQ that might answer the question
- A document that is related to the question

For an example of using the `AMAZON.KendraSearchIntent`, see [Example: Creating a FAQ Bot for an Amazon Kendra Index](#).

If you configure an `AMAZON.KendraSearchIntent` intent for your bot, Amazon Lex calls the intent whenever it can't determine the user utterance for a slot or intent. For example, if your bot is eliciting a response for a slot type called "pizza topping" and the user says "What is a pizza?," Amazon Lex calls the `AMAZON.KendraSearchIntent` to handle the question. If there is no response from Amazon Kendra, the conversation continues as configured in the bot.

When you use both the `AMAZON.KendraSearchIntent` and the `AMAZON.FallbackIntent` in the same bot, Amazon Lex uses the intents as follows:

1. Amazon Lex calls the `AMAZON.KendraSearchIntent`. The intent calls the Amazon Kendra Query operation.
2. If Amazon Kendra returns a response, Amazon Lex displays the result to the user.
3. If there is no response from Amazon Kendra, Amazon Lex re-prompts the user. The next action depends on response from the user.
 - If the response from the user contains an utterance that Amazon Lex recognizes, such as filling a slot value or confirming an intent, the conversation with the user proceeds as configured for the bot.
 - If the response from the user does not contain an utterance that Amazon Lex recognizes, Amazon Lex makes another call to the Query operation.
4. If there is no response after the configured number of retries, Amazon Lex calls the `AMAZON.FallbackIntent` and ends the conversation with the user.

There are three ways to use the `AMAZON.KendraSearchIntent` to make a request to Amazon Kendra:

- Let the search intent make the request for you. Amazon Lex calls Amazon Kendra with the user's utterance as the search string. When you create the intent, you can define a query filter string that limits the number of responses that Amazon Kendra returns. Amazon Lex uses the filter in the query request.
- Add additional query parameters to the request to narrow the search results using your dialog Lambda function. You add a `kendraQueryFilterString` field that contains Amazon Kendra query parameters to the delegate dialog action. When you add query parameters to the request with the Lambda function, they take precedence over the query filter that you defined when you created the intent.
- Create a new query using the dialog Lambda function. You can create a complete Amazon Kendra query request that Amazon Lex sends. You specify the query in the `kendraQueryRequestPayload` field in the delegate dialog action. The `kendraQueryRequestPayload` field takes precedence over the `kendraQueryFilterString` field.

To specify the `queryFilterString` parameter when you create a bot, or to specify the `kendraQueryFilterString` field when you call the delegate action in a dialog Lambda function, you specify a string that is used as the attribute filter for the Amazon Kendra query. If the string isn't a valid attribute filter, you'll get an `InvalidBotConfigException` exception

at runtime. For more information about attribute filters, see [Using document attributes to filter queries](#) in the *Amazon Kendra Developer Guide*.

To have control over the query that Amazon Lex sends to Amazon Kendra, you can specify a query in the `kendraQueryRequestPayload` field in your dialog Lambda function. If the query isn't valid, Amazon Lex returns an `InvalidLambdaResponseException` exception. For more information, see the [Query operation](#) in the *Amazon Kendra Developer Guide*.

For an example of how to use the `AMAZON.KendraSearchIntent`, see [Example: Creating a FAQ Bot for an Amazon Kendra Index](#).

IAM Policy for Amazon Kendra Search

To use the `AMAZON.KendraSearchIntent` intent, you must use a role that provides AWS Identity and Access Management (IAM) policies that enable Amazon Lex to assume a runtime role that has permission to call the Amazon Kendra Query intent. The IAM settings that you use depend on whether you create the `AMAZON.KendraSearchIntent` using the Amazon Lex console, or using an AWS SDK or the AWS Command Line Interface (AWS CLI). When you use the console, you can choose between adding permission to call Amazon Kendra to the Amazon Lex service-linked role or using a role specifically for calling the Amazon Kendra Query operation. When you use the AWS CLI or an SDK to create the intent, you must use a role specifically for calling the Query operation.

Attaching Permissions

You can use the console to attach permissions to access the Amazon Kendra Query operation to the default Amazon Lex service-linked role. When you attach permissions to the service-linked role, you don't have to create and manage a runtime role specifically to connect to the Amazon Kendra index.

The user, role, or group that you use to access the Amazon Lex console must have permissions to manage role policies. Attach the following IAM policy to the console access role. When you grant these permissions, the role has permissions to change the existing service-linked role policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:AttachRolePolicy",
        "iam:PutRolePolicy",

```

```

        "iam:GetRolePolicy"
    ],
    "Resource": "arn:aws:iam::*:role/aws-service-role/lex.amazonaws.com/
AWSServiceRoleForLexBots"
  },
  {
    "Effect": "Allow",
    "Action": "iam:ListRoles",
    "Resource": "*"
  }
]
}

```

Specifying a Role

You can use the console, the AWS CLI, or the API to specify a runtime role to use when calling the Amazon Kendra Query operation.

The user, role, or group that you use to specify the runtime role must have the `iam:PassRole` permission. The following policy defines the permission. You can use the `iam:AssociatedResourceArn` and `iam:PassedToService` condition context keys to further limit the scope of the permissions. For more information, see [IAM and AWS STS Condition Context Keys](#) in the *AWS Identity and Access Management User Guide*.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::account:role/role"
    }
  ]
}

```

The runtime role that Amazon Lex needs to use to call Amazon Kendra must have the `kendra:Query` permissions. When you use an existing IAM role for permission to call the Amazon Kendra Query operation, the role must have the following policy attached.

You can use the IAM console, the IAM API, or the AWS CLI to create a policy and attach it to a role. These instructions use the AWS CLI to create the role and policies.

Note

The following code is formatted for Linux and MacOS. For Windows, replace the Linux line continuation character (\) with a caret (^).

To add Query operation permission to a role

1. Create a document called **KendraQueryPolicy.json** in the current directory, add the following code to it, and save it

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kendra:Query"
      ],
      "Resource": [
        "arn:aws:kendra:region:account:index/index ID"
      ]
    }
  ]
}
```

2. In the AWS CLI, run the following command to create the IAM policy for running the Amazon Kendra Query operation.

```
aws iam create-policy \
  --policy-name query-policy-name \
  --policy-document file://KendraQueryPolicy.json
```

3. Attach the policy to the IAM role that you are using to call the Query operation.

```
aws iam attach-role-policy \
  --policy-arn arn:aws:iam::account-id:policy/query-policy-name
  --role-name role-name
```

You can choose to update the Amazon Lex service-linked role or to use a role that you created when you create the `AMAZON.KendraSearchIntent` for your bot. The following procedure shows how to choose the IAM role to use.

To specify the runtime role for `AMAZON.KendraSearchIntent`

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose the bot that you want to add the `AMAZON.KendraSearchIntent` to.
3. Choose the plus (+) next to **Intents**.
4. In **Add intent**, choose **Search existing intents**.
5. In **Search intents**, enter **`AMAZON.KendraSearchIntent`** and then choose **Add**.
6. In **Copy built-in intent**, enter a name for the intent, such as **`KendraSearchIntent`**, and then choose **Add**.
7. Open the **Amazon Kendra query** section.
8. For **IAM role** choose one of the following options:
 - To update the Amazon Lex service-linked role to enable your bot to query Amazon Kendra indexes, choose **Add Amazon Kendra permissions**.
 - To use a role that has permission to call the Amazon Kendra Query operation, choose **Use an existing role**.

Using Request and Session Attributes as Filters

To filter the response from Amazon Kendra to items related to current conversation, use session and request attributes as filters by adding the `queryFilterString` parameter when you create your bot. You specify a placeholder for the attribute when you create the intent, and then Amazon Lex V2 substitutes a value before it calls Amazon Kendra. For more information about request attributes, see [Setting Request Attributes](#). For more information about session attributes, see [Setting Session Attributes](#).

The following is a example of a `queryFilterString` parameter that uses a string to filter the Amazon Kendra query.

```
"{"equalsTo": {"key": "City", "value": {"stringValue": "Seattle"}}}"
```

The following is an example of a `queryFilterString` parameter that uses a session attribute called "SourceURI" to filter the Amazon Kendra query.

```
"{"equalsTo": {"key": "SourceURI", "value": {"stringValue": "[FileURL]"}}}"
```

The following is an example of a `queryFilterString` parameter that uses a request attribute called "DepartmentName" to filter the Amazon Kendra query.

```
"{"equalsTo": {"key": "Department", "value": {"stringValue": "((DepartmentName))"}}}"
```

The `AMAZON.KendraSearchIntent` filters use the same format as the Amazon Kendra search filters. For more information, see [Using document attributes to filter search results](#) in the *Amazon Kendra developer guide*.

The query filter string used with the `AMAZON.KendraSearchIntent` must use lower-case letters for the first letter of each filter. For example, the following is a valid query filter for the `AMAZON.KendraSearchIntent`.

```
{
  "andAllFilters": [
    {
      "equalsTo": {
        "key": "City",
        "value": {
          "stringValue": "Seattle"
        }
      }
    },
    {
      "equalsTo": {
        "key": "State",
        "value": {
          "stringValue": "Washington"
        }
      }
    }
  ]
}
```

Using the Search Response

Amazon Kendra returns the response to a search in the intent's conclusion statement. The intent must have a conclusion statement unless a fulfillment Lambda function produces a conclusion message.

Amazon Kendra has four types of responses.

- `x-amz-lex:kendra-search-response-question_answer-question-<N>` – The question from a FAQ that matches the search.
- `x-amz-lex:kendra-search-response-question_answer-answer-<N>` – The answer from a FAQ that matches the search.
- `x-amz-lex:kendra-search-response-document-<N>` – An excerpt from a document in the index that is related to the text of the utterance.
- `x-amz-lex:kendra-search-response-document-link-<N>` – The URL of a document in the index that is related to the text of the utterance.
- `x-amz-lex:kendra-search-response-answer-<N>` – An excerpt from a document in the index that answers the question.

The responses are returned in request attributes. There can be up to five responses for each attribute, numbered 1 through 5. For more information about responses, see [Types of response](#) in the *Amazon Kendra Developer Guide*.

The conclusion statement must have one or more message groups. Each message group contains one or more messages. Each message can contain one or more placeholder variables that are replaced by request attributes in the response from Amazon Kendra. There must be at least one message in the message group where all of the variables in the message are replaced by request attribute values in the runtime response, or there must be a message in the group with no placeholder variables. The request attributes are set off with double parentheses ("((" ")). The following message group messages match any response from Amazon Kendra:

- "I found a FAQ question for you: ((x-amz-lex:kendra-search-response-question_answer-question-1)), and the answer is ((x-amz-lex:kendra-search-response-question_answer-answer-1))"
- "I found an excerpt from a helpful document: ((x-amz-lex:kendra-search-response-document-1))"
- "I think the answer to your questions is ((x-amz-lex:kendra-search-response-answer-1))"

Using a Lambda Function to Manage the Request and Response

The `AMAZON.KendraSearchIntent` intent can use your dialog code hook and fulfillment code hook to manage the request to Amazon Kendra and the response. Use the dialog code hook Lambda function when you want to modify the query that you send to Amazon Kendra, and the fulfillment code hook Lambda function when you want to modify the response.

Creating a Query with the Dialog Code Hook

You can use the dialog code hook to create a query to send to Amazon Kendra. Using the dialog code hook is optional. If you don't specify a dialog code hook, Amazon Lex constructs a query from the user utterance and uses the `queryFilterString` that you provided when you configured the intent, if you provided one.

You can use two fields in the dialog code hook response to modify the request to Amazon Kendra:

- `kendraQueryFilterString` – Use this string to specify attribute filters for the Amazon Kendra request. You can filter the query using any of the index fields defined in your index. For the structure of the filter string, see [Using document attributes to filter queries](#) in the *Amazon Kendra Developer Guide*. If the specified filter string isn't valid, you will get an `InvalidLambdaResponseException` exception. The `kendraQueryFilterString` string overrides any query string specified in the `queryFilterString` configured for the intent.
- `kendraQueryRequestPayload` – Use this string to specify an Amazon Kendra query. Your query can use any of the features of Amazon Kendra. If you don't specify a valid query, you get a `InvalidLambdaResponseException` exception. For more information, see [Query](#) in the *Amazon Kendra Developer Guide*.

After you have created the filter or query string, you send the response to Amazon Lex with the `dialogAction` field of the response set to `delegate`. Amazon Lex sends the query to Amazon Kendra and then returns the query response to the fulfillment code hook.

Using the Fulfillment Code Hook for the Response

After Amazon Lex sends a query to Amazon Kendra, the query response is returned to the `AMAZON.KendraSearchIntent` fulfillment Lambda function. The input event to the code hook contains the complete response from Amazon Kendra. The query data is in the same structure as the one returned by the Amazon Kendra Query operation. For more information, see [Query response syntax](#) in the *Amazon Kendra Developer Guide*.

The fulfillment code hook is optional. If one does not exist, or if the code hook doesn't return a message in the response, Amazon Lex uses the conclusion statement for responses.

Example: Creating a FAQ Bot for an Amazon Kendra Index

This example creates an Amazon Lex bot that uses an Amazon Kendra index to provide answers to users' questions. The FAQ bot manages the dialog for the user. It uses the `AMAZON.KendraSearchIntent` intent to query the index and to present the response to the user. To create the bot, you:

1. Create a bot that your customers will interact with to get answers from your bot.
2. Create a custom intent. Your bot requires at least one intent with at least one utterance. This intent enables your bot to build, but is not used otherwise.
3. Add the `KendraSearchIntent` intent to your bot and configure it to work with your Amazon Kendra index.
4. Test the bot by asking questions that are answered by documents stored in your Amazon Kendra index.

Before you can use this example, you need to create an Amazon Kendra index. For more information, see [Getting started with an S3 bucket \(console\)](#) in the *Amazon Kendra Developer Guide*.

To create a FAQ bot

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. In the navigation pane, choose **Bots**.
3. Choose **Create**.
4. Choose **Custom bot**. Configure the bot as follows:
 - **Bot name** – Give the bot a name that indicates its purpose, such as **KendraTestBot**.
 - **Output voice** – Choose **None**.
 - **Session timeout** – Enter **5**.
 - **Sentiment analysis** – Choose **No**.
 - **COPPA** – Choose **No**.
 - **User utterance storage** – Choose **Do not store**.

5. Choose **Create**.

To successfully build a bot, you must create at least one intent with at least one sample utterance. This intent is required to build your Amazon Lex bot, but isn't used for the FAQ response. The utterance for the intent must not apply to any of the questions that your customer asks.

To create the required intent

1. On the **Getting started with your bot** page, choose **Create intent**.
2. For **Add intent**, choose **Create intent**.
3. In the **Create intent** dialog box, give the intent a name, such as **RequiredIntent**.
4. For **Sample utterances**, type an utterance, such as **Required utterance**.
5. Choose **Save intent**.

Now, create the intent to search an Amazon Kendra index and the response messages that it should return.

To create an **AMAZON.KendraSearchIntent** intent and response message

1. In the navigation pane, choose the plus (+) next to **Intents**.
2. For **Add intent**, choose **Search existing intents**.
3. In the **Search intents** box, enter **AMAZON.KendraSearchIntent**, then choose it from the list.
4. For **Copy built-in intent**, give the intent a name, such as **KendraSearchIntent**, and then choose **Add**.
5. In the intent editor, choose **Amazon Kendra query** to open the query options.
6. From the **Amazon Kendra index** menu, choose the index that you want the intent to search.
7. In the **Response** section, add the following three messages:

```
I found a FAQ question for you: ((x-amz-lex:kendra-search-response-question_answer-question-1)) and the answer is ((x-amz-lex:kendra-search-response-question_answer-answer-1)).  
I found an excerpt from a helpful document: ((x-amz-lex:kendra-search-response-document-1)).  
I think the answer to your questions is ((x-amz-lex:kendra-search-response-answer-1)).
```

8. Choose **Save intent**, and then choose **Build** to build the bot.

Finally, use the console test window to test responses from your bot. Your questions should be in the domain that your index supports.

To test your FAQ bot

1. In the console test window, type a question for your index.
2. Verify the answer in the test window's response section.
3. To reset the test window for another question, choose **Clear chat history**.

AMAZON.PauseIntent

Responds to words and phrases that enable the user to pause an interaction with a bot so that they can return to it later. Your Lambda function or application needs to save intent data in session variables, or you need to use the [GetSession](#) operation to retrieve intent data when you resume the current intent.

Common utterances:

- pause
- pause that

AMAZON.RepeatIntent

Responds to words and phrases that enable the user to repeat the previous message. Your application needs to use a Lambda function to save the previous intent information in session variables, or you need to use the [GetSession](#) operation to get the previous intent information.

Common utterances:

- repeat
- say that again
- repeat that

AMAZON.ResumeIntent

Responds to words and phrases that enable the user to resume a previously paused intent. Your Lambda function or application must manage the information required to resume the previous intent.

Common utterances:

- resume
- continue
- keep going

AMAZON.StartOverIntent

Responds to words and phrases that enable the user to stop processing the current intent and start over from the beginning. You can use your Lambda function or the `PutSession` operation to elicit the first slot value again.

Common utterances:

- start over
- restart
- start again

AMAZON.StopIntent

Responds to words and phrases that indicate that the user wants to stop processing the current intent and end the interaction with a bot. Your Lambda function or application should clear any existing attributes and slot type values and then end the interaction.

Common utterances:

- stop
- off
- shut up

Built-in Slot Types

Amazon Lex supports built-in slot types that define how data in the slot is recognized and handled. You can create slots of these types in your intents. This eliminates the need to create enumeration values for commonly used slot data such as date, time, and location. Built-in slot types do not have versions.

Slot Type	Short Description	Supported Locales	
AMAZON.Airport	Recognizes words that represent an airport.	All locales	
AMAZON.AlphaNumeric	Recognizes words made up of letters and numbers.	All locales except Korean (ko-KR)	
AMAZON.City	Recognizes words that represent a city.	All locales	
AMAZON.Country	Recognizes words that represent a country.	All locales	
AMAZON.DATE	Recognizes words that represent a date and converts them to a standard format.	All locales	
AMAZON.DURATION	Recognizes words that represent duration and converts them to a standard format.	All locales	
AMAZON.EmailAddress	Recognizes words that represent an email address and converts them into	All locales	

Slot Type	Short Description	Supported Locales	
	a standard email address.		
AMAZON.FirstName	Recognizes words that represent a first name.	All locales	
AMAZON.LastName	Recognizes words that represent a last name.	All locales	
AMAZON.NUMBER	Recognizes numeric words and converts them into digits.	All locales	
AMAZON.Percentage	Recognizes words that represent a percentage and converts them to a number and a percent sign (%).	All locales	
AMAZON.PhoneNumber	Recognizes words that represent a phone number and converts them into a numeric string.	All locales	
AMAZON.SpeedUnit	Recognizes words that represent a speed unit and converts them into a standard abbreviation.	English (US) (en-US)	

Slot Type	Short Description	Supported Locales	
AMAZON.State	Recognizes words that represent a state.	All locales	
AMAZON.StreetName	Recognizes words that represent a street name.	All locales except English (US) (en-US)	
AMAZON.TIME	Recognizes words that indicate times and converts them into a time format.	All locales	
AMAZON.WeightUnit	Recognizes words that represent a weight unit and converts them into a standard abbreviation	English (US) (en-US)	

Note

For the English (US) (en-US) locale, Amazon Lex supports slot types from the Alexa Skill Kit. For a list of available built-in slot types, see the [Slot Type Reference](#) in the Alexa Skills Kit documentation.

- Amazon Lex doesn't support the `AMAZON.LITERAL` or the `AMAZON.SearchQuery` built-in slot types.

AMAZON.Airport

Provides a list of airports. Examples include:

- John F. Kennedy International Airport
- Melbourne Airport

AMAZON.AlphaNumeric

Recognizes strings made up of letters and numbers, such as **APQ123**.

This slot type is not available in the Korean (ko-KR) locale.

You can use the `AMAZON.AlphaNumeric` slot type for strings that contain:

- Alphabetical characters, such as **ABC**
- Numeric characters, such as **123**
- A combination of alphanumeric characters, such as **ABC123**

You can add a regular expression to the `AMAZON.AlphaNumeric` slot type to validate values entered for the slot. For example, you can use a regular expression to validate:

- United Kingdom or Canadian postal codes
- Driver's license numbers
- Vehicle identification numbers

Use a standard regular expression. Amazon Lex supports the following characters in the regular expression:

- A-Z, a-z
- 0-9

Amazon Lex also supports Unicode characters in regular expressions. The form is `\uUnicode`. Use four digits to represent Unicode characters. For example, `[\u0041-\u005A]` is equivalent to `[A-Z]`.

The following regular expression operators are not supported:

- Infinite repeaters: `*`, `+`, or `{x,}` with no upper bound.
- Wild card (`.`)

The maximum length of the regular expression is 300 characters. The maximum length of a string stored in an `AMAZON.AlphaNumeric` slot type that uses a regular expression is 30 characters.

The following are some example regular expressions.

- Alphanumeric strings, such as **APQ123** or **APQ1**: `[A-Z]{3}[0-9]{1,3}` or a more constrained `[A-DP-T]{3} [1-5]{1,3}`
- US Postal Service Priority Mail International format, such as **CP123456789US**: `CP[0-9]{9}US`
- Bank routing numbers, such as **123456789**: `[0-9]{9}`

To set the regular expression for a slot type, use the console or the [PutSlotType](#) operation. The regular expression is validated when you save the slot type. If the expression isn't valid, Amazon Lex returns an error message.

When you use a regular expression in a slot type, Amazon Lex checks input to slots of that type against the regular expression. If the input matches the expression, the value is accepted for the slot. If the input does not match, Amazon Lex prompts the user to repeat the input.

AMAZON.City

Provides a list of local and world cities. The slot type recognizes common variations of city names. Amazon Lex doesn't convert from a variation to an official name.

Examples:

- New York
- Reykjavik
- Tokyo
- Versailles

AMAZON.Country

The names of countries around the world. Examples:

- Australia
- Germany
- Japan
- United States
- Uruguay

AMAZON.DATE

Converts words that represent dates into a date format.

The date is provided to your intent in ISO-8601 date format. The date that your intent receives in the slot can vary depending on the specific phrase uttered by the user.

- Utterances that map to a specific date, such as "today," "now," or "November twenty-fifth," convert to a complete date: 2020-11-25. This defaults to dates *on or after* the current date.
- Utterances that map to a specific week, such as "this week," or "next week," convert to the date of the first day of the week. In ISO-8601 format, the week starts on Monday and ends on Sunday. For example, if today is 2020-11-25, "next week" converts to 2020-11-30.
- Utterances that map to a month, but not a specific day, such as "next month," convert to the last day of the month. For example, if today is 2020-11-25, "next month" converts to 2020-12-31.
- Utterances that map to a year, but not a specific month or day, such as "next year," convert to the last day of the following year. For example, if today is 2020-11-25, "next year" converts to 2021-12-31.

AMAZON.DURATION

Converts words that indicate durations into a numeric duration.

The duration is resolved to a format based on the [ISO-8601 duration format](#), PnYnMnWnDTnHnMnS. The P indicates that this is a duration, the n is a numeric value, and the capital letter following the n is the specific date or time element. For example, P3D means 3 days. A T is used to indicate that the remaining values represent time elements rather than date elements.

Examples:

- "ten minutes": PT10M
- "five hours": PT5H
- "three days": P3D
- "forty five seconds": PT45S
- "eight weeks": P8W
- "seven years": P7Y
- "five hours ten minutes": PT5H10M

- "two years three hours ten minutes": P2YT3H10M

AMAZON.EmailAddress

Recognizes words that represent an email address provided as username@domain. Addresses can include the following special characters in a user name: underscore (_), hyphen (-), period (.), and the plus sign (+).

AMAZON.FirstName

Commonly used first names. This slot type recognizes both formal names and informal nicknames. The name sent to your intent is the value sent by the user. Amazon Lex doesn't convert from the nick name to the formal name.

For first names that sound alike but are spelled differently, Amazon Lex sends your intent a single common form.

In the English (US) (en-US) locale, use the slot name AMAZON.US_First_Name.

Examples:

- Emily
- John
- Sophie

AMAZON.LastName

Commonly used last names. For names that sound alike that are spelled differently, Amazon Lex sends your intent a single common form.

In the English (US) (en-US) locale, use the slot name AMAZON.US_Last_Name.

Examples:

- Brosky
- Dasher
- Evers
- Parres

- Welt

AMAZON.NUMBER

Converts words or numbers that express a number into digits, including decimal numbers. The following table shows how the AMAZON.NUMBER slot type captures numeric words.

Input	Response
one hundred twenty three point four five	123.45
one hundred twenty three dot four five	123.45
point four two	0.42
point forty two	0.42
232.998	232.998
50	50

AMAZON.Percentage

Converts words and symbols that represent a percentage into a numeric value with a percent sign (%).

If the user enters a number without a percent sign or the word "percent," the slot value is set to the number. The following table shows how the AMAZON.Percentage slot type captures percentages.

Input	Response
50 percent	50%
0.4 percent	0.4%
23.5%	23.5%
twenty five percent	25%

AMAZON.PhoneNumber

Converts the numbers or words that represent a phone number into a string format without punctuation as follows.

Type	Description	Input	Result
International number with leading plus (+) sign	11-digit number with leading plus sign.	+61 7 4445 1061	+61744431061
		+1 (509) 555-1212	+15095551212
International number without leading plus (+) sign	11-digit number without leading plus sign	1 (509) 555-1212	15095551212
		61 7 4445 1061	61744451061
National number	10-digit number without international code	(03) 5115 4444	0351154444
		(509) 555-1212	5095551212
Local number	7-digit phone number without an international code or an area code	555-1212	5551212

AMAZON.SpeedUnit

Converts words that represent speed units into the corresponding abbreviation. For example, "miles per hour" is converted to mph.

This slot type is available only in the English (US) (en-US) locale.

The following examples show how the `AMAZON.SpeedUnit` slot type captures speed units.

Speed unit	Abbreviation
miles per hour, mph, MPH, m/h	mph

Speed unit	Abbreviation
kilometers per hour, km per hour, kmph, KMPH, km/h	kmph
meters per second, mps, MPS, m/s	mps
nautical miles per hour, knots, knot	knot

AMAZON.State

The names of geographical and political regions within countries.

Examples:

- Bavaria
- Fukushima Prefecture
- Pacific Northwest
- Queensland
- Wales

AMAZON.StreetName

The names of streets within a typical street address. This includes just the street name, not the house number.

This slot type isn't available in the English (US) (en-US) locale.

Examples:

- Canberra Avenue
- Front Street
- Market Road

AMAZON.TIME

Converts words that represent times into time values. Includes resolutions for ambiguous times. When a user enters an ambiguous time, Amazon Lex uses the `slotDetails` attribute of a Lambda event to pass resolutions for the ambiguous times to your Lambda function. For example, if your bot prompts the user for a delivery time, the user can respond by saying "10 o'clock." This time is ambiguous. It means either 10:00 AM or 10:00 PM. In this case, the value in the `slots` map is `null`, and the `slotDetails` entity contains the two possible resolutions of the time. Amazon Lex inputs the following into the Lambda function:

```
"slots": {
  "deliveryTime": null
},
"slotDetails": {
  "deliveryTime": {
    "resolutions": [
      {
        "value": "10:00"
      },
      {
        "value": "22:00"
      }
    ]
  }
}
```

When the user responds with an unambiguous time, Amazon Lex sends the time to your Lambda function in the `slots` attribute of the Lambda event and the `slotDetails` attribute is empty. For example, if your user responds to the prompt for a delivery time with "10:00 PM," Amazon Lex inputs the following into the Lambda function:

```
"slots": {
  "deliveryTime": "22:00"
}
```

For more information about the data sent from Amazon Lex to a Lambda function, see [Input Event Format](#).

AMAZON.WeightUnit

Converts words that represent a weight unit into the corresponding abbreviation. For example, "kilogram" is converted to kg.

This slot type is available only in the English (US) (en-US) locale.

The following examples show how the AMAZON.WeightUnit slot type captures weight units:

Weight unit	Abbreviation
kilograms, kilos, kgs, KGS	kg
grams, gms, gm, GMS, g	g
milligrams, mg, mgs	mg
pounds, lbs, LBS	lbs
ounces, oz, OZ	oz
tonne, ton, t	t
kiloton, kt	kt

Custom Slot Types

For each intent, you can specify parameters that indicate the information that the intent needs to fulfill the user's request. These parameters, or slots, have a type. A *slot type* is a list of values that Amazon Lex uses to train the machine learning model to recognize values for a slot. For example, you can define a slot type called "Genres." Each value in the slot type is the name of a genre, "comedy," "adventure," "documentary," etc. You can define a synonym for a slot type value. For example, you can define the synonyms "funny" and "humorous" for the value "comedy."

You can configure the slot type to restrict resolution to the slot values. The slot values will be used as an enumeration and the value entered by the user will be resolved to the slot value only if it is the same as one of the slot values or a synonym. A synonym is resolved to the corresponding slot value. For example, if the user enters "funny" it will resolve to the slot value "comedy."

Alternately, you can configure the slot type to expand the values. Slot values will be used as training data and the slot is resolved to the value provided by the user if it is similar to the slot values and synonyms. This is the default behavior.

Amazon Lex maintains a list of possible resolutions for a slot. Each entry in the list provides a *resolution value* that Amazon Lex recognized as additional possibilities for the slot. A resolution value is the best effort to match the slot value. The list contains up to five values.

When the value entered by the user is a synonym, the first entry in the list of resolution values is the slot type value. For example, if the user enters "funny," the `slots` field contains "funny" and the first entry in the `slotDetails` field is "comedy." You can configure the `valueSelectionStrategy` when you create or update a slot type with the [PutSlotType](#) operation so that the slot value is filled with the first value in the resolution list.

If you are using a Lambda function, the input event to the function includes a resolution list called `slotDetails`. The following example shows the slot and slot details section of the input to a Lambda function:

```
"slots": {
  "MovieGenre": "funny";
},
"slotDetails": {
  "Movie": {
    "resolutions": [
      "value": "comedy"
    ]
  }
}
```

For each slot type, you can define a maximum of 10,000 values and synonyms. Each bot can have a total number of 50,000 slot type values and synonyms. For example, you can have 5 slot types, each with 5,000 values and 5,000 synonyms, or you can have 10 slot types, each with 2,500 values and 2,500 synonyms. If you exceed these limits, you will get a `LimitExceededException` when you call the [PutBot](#) operation.

Slot Obfuscation

Amazon Lex enables you to obfuscate, or hide, the contents of slots so that the content is not visible. To protect sensitive data captured as slot values, you can enable slot obfuscation to mask those values in conversation logs.

When you choose to obfuscate slot values, Amazon Lex replaces the value of the slot with the name of the slot in conversation logs. For a slot called `full_name`, the value of the slot would be obfuscated as follows:

```
Before obfuscation:  
    My name is John Stiles  
After obfuscation:  
    My name is {full_name}
```

If an utterance contains bracket characters (`{}`) Amazon Lex escapes the bracket characters with two back slashes (`\\`). For example, the text `{John Stiles}` is obfuscated as follows:

```
Before obfuscation:  
    My name is {John Stiles}  
After obfuscation:  
    My name is \\{{full_name}}\\
```

Slot values are obfuscated in conversation logs. The slot values are still available in the response from the `PostContent` and `PostText` operations, and the slot values are available to your validation and fulfillment Lambda functions. If you are using slot values in your prompts or responses, those slot values are not obfuscated in conversation logs.

In the first turn of a conversation, Amazon Lex obfuscates slot values if it recognizes a slot and slot value in the utterance. If no slot value is recognized, Amazon Lex does not obfuscate the utterance.

On the second and later turns, Amazon Lex knows the slot to elicit and if the slot value should be obfuscated. If Amazon Lex recognizes the slot value, the value is obfuscated. If Amazon Lex does not recognize a value, the entire utterance is obfuscated. Any slot values in missed utterances won't be obfuscated.

Amazon Lex also doesn't obfuscate slot values that you store in request or session attributes. If you are storing slot values that should be obfuscated as an attribute, you must encrypt or otherwise obfuscate the value.

Amazon Lex doesn't obfuscate the slot value in audio. It does obfuscate the slot value in the audio transcription.

You don't need to obfuscate all of the slots in a bot. You can choose which slots obfuscate using the console or by using the Amazon Lex API. In the console, choose **Slot obfuscation** in the settings for a slot. If you are using the API, set the `obfuscationSetting` field of the slot to `DEFAULT_OBFUSCATION` when you call the [PutIntent](#) operation.

Sentiment Analysis

You can use sentiment analysis to determine the sentiments expressed in a user utterance. With the sentiment information you can manage conversation flow or perform post-call analysis. For example, if the user sentiment is negative you can create a flow to hand over a conversation to a human agent.

Amazon Lex integrates with Amazon Comprehend to detect user sentiment. The response from Amazon Comprehend indicates whether the overall sentiment of the text is positive, neutral, negative, or mixed. The response contains the most likely sentiment for the user utterance and the scores for each of the sentiment categories. The score represents the likelihood that the sentiment was correctly detected.

You enable sentiment analysis for a bot using the console or by using the Amazon Lex API. On the Amazon Lex console, choose the **Settings** tab for your bot, then set the **Sentiment Analysis** option to **Yes**. If you are using the API, call the [PutBot](#) operation with the `detectSentiment` field set to `true`.

When sentiment analysis is enabled, the response from the [PostContent](#) and [PostText](#) operations return a field called `sentimentResponse` in the bot response with other metadata. The `sentimentResponse` field has two fields, `SentimentLabel` and `SentimentScore`, that contain the result of the sentiment analysis. If you are using a Lambda function, the `sentimentResponse` field is included in the event data sent to your function.

The following is an example of the `sentimentResponse` field returned as part of the `PostText` or `PostContent` response. The `SentimentScore` field is a string that contains the scores for the response.

```
{
  "SentimentScore":
    "{
```

```
Mixed: 0.030585512690246105,  
Positive: 0.94992071056365967,  
Neutral: 0.0141543131828308,  
Negative: 0.00893945890665054  
}],  
"SentimentLabel": "POSITIVE"  
}
```

Amazon Lex calls Amazon Comprehend on your behalf to determine the sentiment in every utterance processed by the bot. By enabling sentiment analysis, you agree to the service terms and agreements for Amazon Comprehend. For more information about pricing for Amazon Comprehend, see [Amazon Comprehend Pricing](#).

For more information about how Amazon Comprehend sentiment analysis works, see [Determine the Sentiment](#) in the *Amazon Comprehend Developer Guide*.

Tagging Your Amazon Lex Resources

To help you manage your Amazon Lex bots, bot aliases, and bot channels, you can assign metadata to each resource as *tags*. A tag is a label that you assign to an AWS resource. Each tag consists of a key and a value.

Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or application. Tags help you to:

- Identify and organize your AWS resources. Many AWS resources support tagging, so you can assign the same tag to resources in different services to indicate that the resources are related. For example, you can tag a bot and the Lambda functions that it uses with the same tag.
- Allocate costs. You activate tags on the AWS Billing and Cost Management dashboard. AWS uses the tags to categorize your costs and deliver a monthly cost allocation report to you. For Amazon Lex, you can allocate costs for each alias using tags specific to the alias, except for the \$LATEST alias. You allocate costs for the \$LATEST alias using tags for your Amazon Lex bot. For more information, see [Use Cost Allocation Tags](#) in the *AWS Billing and Cost Management User Guide*.
- Control access to your resources. You can use tags to Amazon Lex to create policies to control access to Amazon Lex resources. These policies can be attached to an IAM role or user to enable tag-based access control. For more information, see [ABAC with Amazon Lex](#). To view an example identity-based policy for limiting access to a resource based on the tags on that resource, see [Use a Tag to Access a Resource](#).

You can work with tags using the AWS Management Console, the AWS Command Line Interface, or the Amazon Lex API.

Tagging Your Resources

If you are using the Amazon Lex console, you can tag resources when you create them, or you can add the tags later. You can also use the console to update or remove existing tags.

If you are using the AWS CLI or the Amazon Lex API, you use the following operations to manage tags for your resources:

- [ListTagsForResource](#) – view the tags associated with a resource.
- [PutBot](#) and [PutBotAlias](#) – apply tags when you create a bot or a bot alias.
- [TagResource](#) – add and modify tags on an existing resource.
- [UntagResource](#) – remove tags from a resource.

The following resources in Amazon Lex support tagging:

- Bots - use an Amazon Resource Name (ARN) like the following:
 - `arn:${partition}:lex:${region}:${account}:bot:${bot-name}`
- Bot aliases - use an ARN like the following:
 - `arn:${partition}:lex:${region}:${account}:bot:${bot-name}:${bot-alias}`
- Bot channels - use an ARN like the following:
 - `arn:${partition}:lex:${region}:${account}:bot-channel:${bot-name}:${bot-alias}:${channel-name}`

Tag Restrictions

The following basic restrictions apply to tags on Amazon Lex resources:

- Maximum number of tags - 50
- Maximum key length – 128 characters
- Maximum value length – 256 characters
- Valid characters for key and value – a–z, A–Z, 0–9, space, and the following characters: `_ . : / = + -` and `@`

- Keys and values are case sensitive.
- Don't use `aws :` as a prefix for keys; it's reserved for AWS use.

Tagging Resources (Console)

You can use the console to manage tags on a bot, a bot alias, or a bot channel resource. You can add tags when you create a resource, or you can add, modify, or remove tags from existing resources.

To add a tag when you create a bot

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose **Create** to create a new bot.
3. At the bottom of the **Create your bot** page, choose **Tags**.
4. Choose **Add tag** and add one or more tags to the bot. You can add up to 50 tags.

To add a tag when you create a bot alias

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose the bot that you want to add the bot alias to.
3. Choose **Settings**.
4. Add the alias name, choose the bot version, and then choose **Add tags**.
5. Choose **Add tag** and add one or more tags to the bot alias. You can add up to 50 tags.

To add a tag when you create a bot channel

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose the bot that you want to add the bot channel to.
3. Choose **Channels** and then choose the channel that you want to add.
4. Add the details for the bot channel, and then choose **Tags**.
5. Choose **Add tag** and add one or more tags to the bot channel. You can add up to 50 tags.

To add a tag when you import a bot

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose **Actions** and then choose **Import**.
3. Choose the zip file for importing the bot.
4. Choose **Tags**, then choose **Add tag** to add one or more tags to the bot. You can add up to 50 tags.

To add, remove, or modify a tag on an existing bot

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. From the left menu, choose **Bots** and then choose the bot that you want to modify.
3. Choose **Settings** and then from the left menu choose **General**.
4. Choose **Tags** and then add, modify, or remove tags for the bot.

To add, remove, or modify a tag on a bot alias

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. From the left menu, choose **Bots** and then choose the bot that you want to modify.
3. Choose **Settings** and then from the left menu choose **Aliases**.
4. Choose **Manage tags** for the alias that you want to modify, and then add, modify, or remove tags for the bot alias.

To add, remove, or modify a tag on an existing bot channel

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. From the left menu, choose **Bots** and then choose the bot that you want to modify.
3. Choose **Channels**.
4. Choose **Tags** and then add, modify, or remove tags for the bot channel.

Tagging Resources (AWS CLI)

You can use the AWS CLI to manage tags on a bot, a bot alias, or a bot channel resource. You can add tags when you create a bot or a bot alias, or you can add, modify, or remove tags from a bot, a bot alias, or a bot channel.

All of the examples are formatted for Linux and macOS. To use the command in Windows, replace the Linux continuation character (\) with a caret (^).

To add a tag when you create a bot

- The following abbreviated `put-bot` AWS CLI command shows the parameters that you must use to add a tag when you create a bot. To actually create a bot, you must supply other parameters. For more information, see [Step 4: Getting Started \(AWS CLI\)](#).

```
aws lex-models put-bot \  
  --tags '[{"key": "key1", "value": "value1"}, \  
          {"key": "key2", "value": "value2"}]'
```

To add a tag when you create a bot alias

- The following abbreviated `put-bot-alias` AWS CLI command shows the parameters that you must use to add a tag when you create a bot alias. To actually create a bot alias, you must supply other parameters. For more information, see [Exercise 5: Create an Alias \(AWS CLI\)](#).

```
aws lex-models put-bot \  
  --tags '[{"key": "key1", "value": "value1"}, \  
          {"key": "key2", "value": "value2"}]'
```

To list tags on a resource

- Use the `list-tags-for-resource` AWS CLI command to show the resources associated with a bot, bot alias, bot channel.

```
aws lex-models list-tags-for-resource \  
  --resource-arn bot, bot alias, or bot channel ARN
```

To add or modify tags on a resource

- Use the `tag-resource` AWS CLI command to add or modify a bot, bot alias, or bot channel.

```
aws lex-models tag-resource \  
  --resource-arn bot, bot alias, or bot channel ARN \  
  --tags '[{"key": "key1", "value": "value1"}, \  
          {"key": "key2", "value": "value2"}]'
```

To remove tags from a resource

- Use the `untag-resource` AWS CLI command to remove tags from a bot, bot alias, or bot channel.

```
aws lex-models untag-resource \  
  --resource-arn bot, bot alias, or bot channel ARN \  
  --tag-keys '["key1", "key2"]'
```

Getting Started with Amazon Lex

Amazon Lex provides API operations that you can integrate with your existing applications. For a list of supported operations, see the [API Reference](#). You can use any of the following options:

- **AWS SDK** — When using the SDKs your requests to Amazon Lex are automatically signed and authenticated using the credentials that you provide. This is the recommended choice for building your applications.
- **AWS CLI** — You can use the AWS CLI to access any Amazon Lex feature without having to write any code.
- **AWS Console** — The console is the easiest way to get started testing and using Amazon Lex

If you are new to Amazon Lex, we recommend that you read [Amazon Lex: How It Works](#). first.

Topics

- [Step 1: Set Up an AWS Account and Create an Administrator User](#)
- [Step 2: Set Up the AWS Command Line Interface](#)
- [Step 3: Getting Started \(Console\)](#)
- [Step 4: Getting Started \(AWS CLI\)](#)

Step 1: Set Up an AWS Account and Create an Administrator User

Before you use Amazon Lex for the first time, complete the following tasks:

1. [Sign Up for AWS](#)
2. [Create a user](#)

Sign Up for AWS

If you already have an AWS account, skip this task.

When you sign up for Amazon Web Services (AWS), your AWS account is automatically signed up for all services in AWS, including Amazon Lex. You are charged only for the services that you use.

With Amazon Lex, you pay only for the resources that you use. If you are a new AWS customer, you can get started with Amazon Lex for free. For more information, see [AWS Free Usage Tier](#).

If you already have an AWS account, skip to the next task. If you don't have an AWS account, use the following procedure to create one.

To create an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

Write down your AWS account ID because you'll need it for the next task.

Create a user

Services in AWS, such as Amazon Lex, require that you provide credentials when you access them so that the service can determine whether you have permissions to access the resources owned by that service. The console requires your password. However, we don't recommend that you access AWS using the credentials for your AWS account. Instead, we recommend that you:

- Use AWS Identity and Access Management (IAM) to create a user
- Add the user to an IAM group with administrative permissions
- Grant administrative permissions to the user that you created.

You can then access AWS using a special URL and the user's credentials.

The Getting Started exercises in this guide assume that you have a user (`adminuser`) with administrator privileges. Follow the procedure to create `adminuser` in your account.

To create an administrator user and sign in to the console

1. Create an administrator user called `adminuser` in your AWS account. For instructions, see [Creating Your First User and Administrators Group](#) in the *IAM User Guide*.
2. As a user, you can sign in to the AWS Management Console using a special URL. For more information, [How Users Sign In to Your Account](#) in the *IAM User Guide*.

For more information about IAM, see the following:

- [AWS Identity and Access Management \(IAM\)](#)
- [Getting started](#)
- [IAM User Guide](#)

Next Step

[Step 2: Set Up the AWS Command Line Interface](#)

Step 2: Set Up the AWS Command Line Interface

If you prefer to use Amazon Lex with the AWS Command Line Interface (AWS CLI), download and configure it.

Important

You don't need the AWS CLI to perform the steps in the Getting Started exercises. However, some of the later exercises in this guide use the AWS CLI. If you prefer to start by using the console, skip this step and go to [Step 3: Getting Started \(Console\)](#). Later, when you need the AWS CLI, return here to set it up.

To set up the AWS CLI

1. Download and configure the AWS CLI. For instructions, see the following topics in the *AWS Command Line Interface User Guide*:
 - [Getting Set Up with the AWS Command Line Interface](#)
 - [Configuring the AWS Command Line Interface](#)

2. Add a named profile for the administrator user to the end of the AWS CLI config file. You use this profile when executing AWS CLI commands. For more information about named profiles, see [Named Profiles](#) in the *AWS Command Line Interface User Guide*.

```
[profile adminuser]
aws_access_key_id = adminuser access key ID
aws_secret_access_key = adminuser secret access key
region = aws-region
```

For a list of available AWS Regions, see [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

3. Verify the setup by typing the Help command at the command prompt:

```
aws help
```

[Step 3: Getting Started \(Console\)](#)

Step 3: Getting Started (Console)

The easiest way to learn how to use Amazon Lex is by using the console. To get you started, we created the following exercises, all of which use the console:

- Exercise 1 — Create an Amazon Lex bot using a blueprint, a predefined bot that provides all of the necessary bot configuration. You do only a minimum of work to test the end-to-end setup.

In addition, you use the Lambda function blueprint, provided by AWS Lambda, to create a Lambda function. The function is a code hook that uses predefined code that is compatible with your bot.

- Exercise 2 — Create a custom bot by manually creating and configuring a bot. You also create a Lambda function as a code hook. Sample code is provided.
- Exercise 3 — Publish a bot, and then create a new version of it. As part of this exercise you create an alias that points to the bot version.

Topics

- [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#)

- [Exercise 2: Create a Custom Amazon Lex Bot](#)
- [Exercise 3: Publish a Version and Create an Alias](#)

Exercise 1: Create an Amazon Lex Bot Using a Blueprint (Console)

In this exercise, you do the following:

- Create your first Amazon Lex bot, and test it in the Amazon Lex console.

For this exercise, you use the **OrderFlowers** blueprint. For information about blueprints, see [Amazon Lex and AWS Lambda Blueprints](#).

- Create an AWS Lambda function and test it in the Lambda console. While processing a request, your bot calls this Lambda function. For this exercise, you use a Lambda blueprint (**lex-order-flowers-python**) provided in the AWS Lambda console to create your Lambda function. The blueprint code illustrates how you can use the same Lambda function to perform initialization and validation, and to fulfill the `OrderFlowers` intent.
- Update the bot to add the Lambda function as the code hook to fulfill the intent. Test the end-to-end experience.

The following sections explain what the blueprints do.

Amazon Lex Bot: Blueprint Overview

You use the **OrderFlowers** blueprint to create an Amazon Lex bot. For more information about the structure of a bot, see [Amazon Lex: How It Works](#). The bot is preconfigured as follows:

- **Intent** – `OrderFlowers`
- **Slot types** – One custom slot type called `FlowerTypes` with enumeration values: `roses`, `lilies`, and `tulips`.
- **Slots** – The intent requires the following information (that is, slots) before the bot can fulfill the intent.
 - `PickupTime` (`AMAZON.TIME` built-in type)
 - `FlowerType` (`FlowerTypes` custom type)

- **PickupDate** (AMAZON.DATE built-in type)
- **Utterance** – The following sample utterances indicate the user's intent:
 - "I would like to pick up flowers."
 - "I would like to order some flowers."
- **Prompts** – After the bot identifies the intent, it uses the following prompts to fill the slots:
 - Prompt for the `FlowerType` slot – "What type of flowers would you like to order?"
 - Prompt for the `PickupDate` slot – "What day do you want the {FlowerType} to be picked up?"
 - Prompt for the `PickupTime` slot – "At what time do you want the {FlowerType} to be picked up?"
 - Confirmation statement – "Okay, your {FlowerType} will be ready for pickup by {PickupTime} on {PickupDate}. Does this sound okay?"

AWS Lambda Function: Blueprint Summary

The Lambda function in this exercise performs both initialization and validation and fulfillment tasks. Therefore, after creating the Lambda function, you update the intent configuration by specifying the same Lambda function as a code hook to handle both the initialization and validation and fulfillment tasks.

- As an initialization and validation code hook, the Lambda function performs basic validation. For example, if the user provides a time for pickup that is outside of normal business hours, the Lambda function directs Amazon Lex to re-prompt the user for the time.
- As part of the fulfillment code hook, the Lambda function returns a summary message indicating that the flower order has been placed (that is, the intent is fulfilled).

Next Step

[Step 1: Create an Amazon Lex Bot \(Console\)](#)

Step 1: Create an Amazon Lex Bot (Console)

For this exercise, create a bot for ordering flowers, called `OrderFlowersBot`.


To create an Amazon Lex bot (console)

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. If this is your first bot, choose **Get Started**; otherwise, on the **Bots** page, choose **Create**.
3. On the **Create your Lex bot** page, provide the following information, and then choose **Create**.
 - Choose the **OrderFlowers** blueprint.
 - Leave the default bot name (OrderFlowers).
 - For **COPPA**, choose **No**.
 - For **User utterance storage**, choose the appropriate response.
4. Choose **Create**. The console makes the necessary requests to Amazon Lex to save the configuration. The console then displays the bot editor window.
5. Wait for confirmation that your bot was built.
6. Test the bot.

 Note

You can test the bot by typing text into the test window, or, for compatible browsers, by choosing the microphone button in the test window and speaking.

Use the following example text to engage in conversation with the bot to order flowers:



From this input, the bot infers the `OrderFlowers` intent and prompts for slot data. When you provide all of the required slot data, the bot fulfills the intent (`OrderFlowers`) by returning all of the information to the client application (in this case, the console). The console shows the information in the test window.

Specifically:

- In the statement "What day do you want the roses to be picked up?,"the term "roses" appears because the prompt for the `pickupDate` slot is configured using substitutions, `{FlowerType}`. Verify this in the console.
- The "Okay, your roses will be ready..." statement is the confirmation prompt that you configured.
- The last statement ("`FlowerType:roses . . .`") is just the slot data that is returned to the client, in this case, in the test window. In the next exercise, you use a Lambda function to fulfill the intent, in which case you get a message indicating that the order is fulfilled.

Next Step

[Step 2 \(Optional\): Review the Details of Information Flow \(Console\)](#)

Step 2 (Optional): Review the Details of Information Flow (Console)

This section explains the flow of information between a client and Amazon Lex for each user input in our example conversation.

The example uses the console test window for the conversation with the bot.

To open the Amazon Lex test window

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose the bot to test.
3. From the right side of the console, choose **Test chatbot**.

To see the flow of information for spoken or typed content, choose the appropriate topic.

Topics

- [Step 2a \(Optional\): Review the Details of the Spoken Information Flow \(Console\)](#)
- [Step 2b \(Optional\): Review the Details of the Typed Information Flow \(Console\)](#)

Step 2a (Optional): Review the Details of the Spoken Information Flow (Console)

This section explains the flow of information between the client and Amazon Lex when the client uses speech to send requests. For more information, see [PostContent](#).

1. The user says: I would like to order some flowers.
 - a. The client (console) sends the following [PostContent](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/
content HTTP/1.1
x-amz-lex-session-attributes: "e30="
Content-Type: "audio/x-l16; sample-rate=16000; channel-count=1"
Accept: "audio/mpeg"
```

Request body

input stream

Both the request URI and the body provide information to Amazon Lex:

- Request URI – Provides the bot name (*OrderFlowers*), bot alias (*\$LATEST*), and the user name (a random string that identifies the user). *content* indicates that this is a *PostContent* API request (not a *PostText* request).
- Request headers
 - *x-amz-lex-session-attributes* – The base64-encoded value represents "{}". When the client makes the first request, there are no session attributes.
 - *Content-Type* – Reflects the audio format.
- Request body – The user input audio stream ("I would like to order some flowers.").

Note

If the user chooses to send text ("I would like to order some flowers") to the *PostContent* API instead of speaking, the request body is the user input. The *Content-Type* header is set accordingly:

```
POST /bot/OrderFlowers/alias/$LATEST/
user/4o9wwdhx6nlheferh6a73fujd3118f5w/content HTTP/1.1
x-amz-lex-session-attributes: "e30="
Content-Type: "text/plain; charset=utf-8"
Accept: accept
```

Request body

input stream

- b. From the input stream, Amazon Lex detects the intent (`OrderFlowers`). It then chooses one of the intent's slots (in this case, the `FlowerType`) and one of its value elicitation prompts, and then sends a response with the following headers:

```
x-amz-lex-dialog-state:ElicitSlot
x-amz-lex-input-transcript:I would like to order some flowers.
x-amz-lex-intent-name:OrderFlowers
x-amz-lex-message:What type of flowers would you like to order?
x-amz-lex-session-attributes:e30=
x-amz-lex-slot-to-elicite:FlowerType
x-amz-lex-
slots:eyJQaWNrdXBuW1lIjpudWxsLCJGbG93ZXJueXB1IjpudWxsLCJQaWNrdXBeyXR1IjpudWxsQ==
```

The header values provide the following information:

- `x-amz-lex-input-transcript` – Provides the transcript of the audio (user input) from the request
- `x-amz-lex-message` – Provides the transcript of the audio Amazon Lex returned in the response
- `x-amz-lex-slots` – The base64 encoded version of the slots and values:

```
{"PickupTime":null,"FlowerType":null,"PickupDate":null}
```

- `x-amz-lex-session-attributes` – The base64-encoded version of the session attributes ({}).

The client plays the audio in the response body.

2. The user says: roses

- a. The client (console) sends the following [PostContent](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/
content HTTP/1.1
x-amz-lex-session-attributes: "e30="
Content-Type: "audio/x-l16; sample-rate=16000; channel-count=1"
Accept: "audio/mpeg"
```

```
Request body


```

The request body is the user input audio stream (roses). The `sessionAttributes` remains empty.

- b. Amazon Lex interprets the input stream in the context of the current intent (it remembers that it had asked this user for information pertaining to the `FlowerType` slot). Amazon Lex first updates the slot value for the current intent. It then chooses another slot (`PickupDate`), along with one of its prompt messages (When do you want to pick up the roses?), and returns a response with the following headers:

```
x-amz-lex-dialog-state:ElicitSlot
x-amz-lex-input-transcript:roses
x-amz-lex-intent-name:OrderFlowers
x-amz-lex-message:When do you want to pick up the roses?
x-amz-lex-session-attributes:e30=
x-amz-lex-slot-to-elicite:PickupDate
x-amz-lex-
slots:eyJQaWNrdXBuaW1lIjpuWxsLCJGbG93ZXJueXB1Ijoicm9zaSdzIiwUGlja3VwRGF0ZSI6bnVsbH0=
```

The header values provide the following information:

- `x-amz-lex-slots` – The base64-encoded version of the slots and values:

```
{"PickupTime":null,"FlowerType":"roses","PickupDate":null}
```

- `x-amz-lex-session-attributes` – The base64-encoded version of the session attributes ({}).

The client plays the audio in the response body.

3. The user says: tomorrow

- a. The client (console) sends the following [PostContent](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/
content HTTP/1.1
x-amz-lex-session-attributes: "e30="
Content-Type: "audio/x-l16; sample-rate=16000; channel-count=1"
```

```
Accept: "audio/mpeg"
```

Request body

```
input stream ("tomorrow")
```

The request body is the user input audio stream ("tomorrow"). The sessionAttributes remains empty.

- b. Amazon Lex interprets the input stream in the context of the current intent (it remembers that it had asked this user for information pertaining to the PickupDate slot). Amazon Lex updates the slot (PickupDate) value for the current intent. It then chooses another slot to elicit value for (PickupTime) and one of the value elicitation prompts (When do you want to pick up the roses on 2017-03-18?), and returns a response with the following headers:

```
x-amz-lex-dialog-state:ElicitSlot
x-amz-lex-input-transcript:tomorrow
x-amz-lex-intent-name:OrderFlowers
x-amz-lex-message:When do you want to pick up the roses on 2017-03-18?
x-amz-lex-session-attributes:e30=
x-amz-lex-slot-to-elicite:PickupTime
x-amz-lex-
slots:eyJQaWNrdXBuaW11IjpuYWxsLCJGbg93ZXJ0eXB1Ijoicm9zaSdzIiwuUG1ja3VwRGF0ZSI6IjIwMTctM
x-amzn-RequestId:3a205b70-0b69-11e7-b447-eb69face3e6f
```

The header values provide the following information:

- x-amz-lex-slots – The base64-encoded version of the slots and values:

```
{"PickupTime":null,"FlowerType":"roses","PickupDate":"2017-03-18"}
```

- x-amz-lex-session-attributes – The base64-encoded version of the session attributes ({})

The client plays the audio in the response body.

4. The user says: 6 pm

- a. The client (console) sends the following [PostContent](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6n1heferh6a73fujd3118f5w/
content HTTP/1.1
x-amz-lex-session-attributes: "e30="
Content-Type: "text/plain; charset=utf-8"
Accept: "audio/mpeg"
```

Request body

```
input stream ("6 pm")
```

The request body is the user input audio stream ("6 pm"). The sessionAttributes remains empty.

- b. Amazon Lex interprets the input stream in the context of the current intent (it remembers that it had asked this user for information pertaining to the `PickupTime` slot). It first updates the slot value for the current intent.

Now Amazon Lex detects that it has information for all of the slots. However, the `OrderFlowers` intent is configured with a confirmation message. Therefore, Amazon Lex needs an explicit confirmation from the user before it can proceed to fulfill the intent. It sends a response with the following headers requesting confirmation before ordering the flowers:

```
x-amz-lex-dialog-state:ConfirmIntent
x-amz-lex-input-transcript:six p. m.
x-amz-lex-intent-name:OrderFlowers
x-amz-lex-message:Okay, your roses will be ready for pickup by 18:00 on
  2017-03-18. Does this sound okay?
x-amz-lex-session-attributes:e30=
x-amz-lex-
slots:eyJQaWNrdXBuaW1lIjoimTG6MDAiLCJGbG93ZXJlIjoicm9zaSdzIiwiaUglja3VwRGF0ZSI6IjIwM
x-amzn-RequestId:083ca360-0b6a-11e7-b447-eb69face3e6f
```

The header values provide the following information:

- `x-amz-lex-slots` – The base64-encoded version of the slots and values:

```
{"PickupTime":"18:00","FlowerType":"roses","PickupDate":"2017-03-18"}
```


- `x-amz-lex-session-attributes` – The base64-encoded version of the session attributes ({}).

The client plays the audio in the response body.

5. The user says: Yes

- a. The client (console) sends the following [PostContent](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/
content HTTP/1.1
x-amz-lex-session-attributes: "e30="
Content-Type: "audio/x-l16; sample-rate=16000; channel-count=1"
Accept: "audio/mpeg"
```

Request body
input stream ("Yes")

The request body is the user input audio stream ("Yes"). The `sessionAttributes` remains empty.

- b. Amazon Lex interprets the input stream and understands that the user wants to proceed with the order. The `OrderFlowers` intent is configured with `ReturnIntent` as the fulfillment activity. This directs Amazon Lex to return all of the intent data to the client. Amazon Lex returns a response with the following:

```
x-amz-lex-dialog-state:ReadyForFulfillment
x-amz-lex-input-transcript:yes
x-amz-lex-intent-name:OrderFlowers
x-amz-lex-session-attributes:e30=
x-amz-lex-
slots:eyJQaWNrdXBuaw1lIjoimTg6MDAiLCJGbG93ZXJueXB1Ijoicm9zaSdzIiwUGlja3VwRGF0ZSI6IjIwMj
```

The `x-amz-lex-dialog-state` response header is set to `ReadyForFulfillment`. The client can then fulfill the intent.

6. Now, retest the bot. To establish a new (user) context, choose the **Clear** link in the console. Provide data for the `OrderFlowers` intent, and include some invalid data. For example:

- Jasmine as the flower type (it is not one of the supported flower types)
- Yesterday as the day when you want to pick up the flowers

Notice that the bot accepts these values because you don't have any code to initialize and validate the user data. In the next section, you add a Lambda function to do this. Note the following about the Lambda function:

- It validates slot data after every user input. It fulfills the intent at the end. That is, the bot processes the flower order and returns a message to the user instead of simply returning slot data to the client. For more information, see [Using Lambda Functions](#).
- It also sets the session attributes. For more information about session attributes, see [PostText](#).

After you complete the Getting Started section, you can do the additional exercises ([Additional Examples: Creating Amazon Lex Bots](#)). [Book Trip](#) uses session attributes to share cross-intent information to engage in a dynamic conversation with the user.

Next Step

[Step 3: Create a Lambda Function \(Console\)](#)

Step 2b (Optional): Review the Details of the Typed Information Flow (Console)

This section explains flow of information between client and Amazon Lex in which the client uses the `PostText` API to send requests. For more information, see [PostText](#).

1. User types: I would like to order some flowers
 - a. The client (console) sends the following [PostText](#) request to Amazon Lex:


```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/text
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"

{
  "inputText": "I would like to order some flowers",
  "sessionAttributes": {}
}
```

Both the request URI and the body provide information to Amazon Lex:

- Request URI – Provides bot name (`OrderFlowers`), bot alias (`$LATEST`), and user name (a random string identifying the user). The trailing text indicates that it is a `PostText` API request (and not `PostContent`).
 - Request body – Includes the user input (`inputText`) and empty `sessionAttributes`. When the client makes the first request, there are no session attributes. The Lambda function initiates them later.
- b. From the `inputText`, Amazon Lex detects the intent (`OrderFlowers`). This intent does not have any code hooks (that is, the Lambda functions) for initialization and validation of user input or fulfillment.

Amazon Lex chooses one of the intent's slots (`FlowerType`) to elicit the value. It also selects one of the value-elicitation prompts for the slot (all part of the intent configuration), and then sends the following response back to the client. The console displays the message in the response to the user.



The client displays the message in the response.

2. User types: roses

- a. The client (console) sends the following [PostText](#) request to Amazon Lex:


```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/text
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"

{
  "inputText": "roses",
  "sessionAttributes": {}
}
```

The `inputText` in the request body provides user input. The `sessionAttributes` remains empty.

- b. Amazon Lex first interprets the `inputText` in the context of the current intent—the service remembers that it had asked the specific user for information about the `FlowerType` slot. Amazon Lex first updates the slot value for the current intent and chooses another slot (`PickupDate`) along with one of its prompt messages—What day do you want the roses to be picked up?— for the slot.

Then, Amazon Lex returns the following response:



The client displays the message in the response.

3. User types: tomorrow

- a. The client (console) sends the following [PostText](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/text
```


```
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"

{
  "inputText": "tomorrow",
  "sessionAttributes": {}
}
```

The `inputText` in the request body provides user input. The `sessionAttributes` remains empty.

- b. Amazon Lex first interprets the `inputText` in the context of the current intent—the service remembers that it had asked the specific user for information about the `PickupDate` slot. Amazon Lex updates the slot (`PickupDate`) value for the current intent. It chooses another slot to elicit value for (`PickupTime`). It returns one of the value-elicitation prompts—`Deliver the roses at what time on 2017-01-05?`—to the client.

Amazon Lex then returns the following response:



The client displays the message in the response.

4. User types: 6 pm

- a. The client (console) sends the following [PostText](#) request to Amazon Lex:


```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/text
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"
```

```
{
  "inputText": "6 pm",
  "sessionAttributes": {}
}
```

The `inputText` in the request body provides user input. The `sessionAttributes` remains empty.

- b. Amazon Lex first interprets the `inputText` in the context of the current intent—the service remembers that it had asked the specific user for information about the `PickupTime` slot. Amazon Lex first updates the slot value for the current intent. Now Amazon Lex detects that it has information for all the slots.

The `OrderFlowers` intent is configured with a confirmation message. Therefore, Amazon Lex needs an explicit confirmation from the user before it can proceed to fulfill the intent. Amazon Lex sends the following message to the client requesting confirmation before ordering the flowers:



The client displays the message in the response.

5. User types: Yes


- a. The client (console) sends the following [PostText](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/4o9wwdhx6nlheferh6a73fujd3118f5w/text
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"
```

```
{
  "inputText": "Yes",
  "sessionAttributes": {}
}
```

The `inputText` in the request body provides user input. The `sessionAttributes` remains empty.

- b. Amazon Lex interprets the `inputText` in the context of confirming the current intent. It understands that the user want to proceed with the order. The `OrderFlowers` intent is configured with `ReturnIntent` as the fulfillment activity (there is no Lambda function to fulfill the intent). Therefore, Amazon Lex returns the following slot data to the client.



Amazon Lex set the `dialogState` to `ReadyForFulfillment`. The client can then fulfill the intent.

6. Now test the bot again. To do that, you must choose the **Clear** link in the console to establish a new (user) context. Now as you provide data for the order flowers intent, try to provide invalid data. For example:
- Jasmine as the flower type (it is not one of the supported flower types).
 - Yesterday as the day when you want to pick up the flowers.

Notice that the bot accepts these values because you don't have any code to initialize/validate user data. In the next section, you add a Lambda function to do this. Note the following about the Lambda function:

- The Lambda function validates slot data after every user input. It fulfills the intent at the end. That is, the bot processes the flowers order and returns a message to the user instead of simply returning slot data to the client. For more information, see [Using Lambda Functions](#).
- The Lambda function also sets the session attributes. For more information about session attributes, see [PostText](#).

After you complete the Getting Started section, you can do the additional exercises ([Additional Examples: Creating Amazon Lex Bots](#)). [Book Trip](#) uses session attributes to share cross-intent information to engage in a dynamic conversation with the user.

Next Step

[Step 3: Create a Lambda Function \(Console\)](#)

Step 3: Create a Lambda Function (Console)

Create a Lambda function (using the **lex-order-flowers-python** blueprint) and perform test invocation using sample event data in the AWS Lambda console.

You return to the Amazon Lex console and add the Lambda function as the code hook to fulfill the `OrderFlowers` intent in the `OrderFlowersBot` that you created in the preceding section.

To create the Lambda function (console)

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. Choose **Create function**.
3. On the **Create function** page, choose **Use a blueprint**. Type **lex-** in the filter text box and then press Enter to find the blueprint, choose the `lex-order-flowers-python` blueprint.

Lambda function blueprints are provided in both Node.js and Python. For this exercise, use the Python-based blueprint.

4. On the **Basic information** page, do the following.
 - Type a Lambda function name (`OrderFlowersCodeHook`).
 - For the execution role, choose **Create a new role with basic Lambda permissions**.

- Leave the other default values.
- 5. Choose **Create function**.
- 6. If you are using a locale other than English (US) (en-US), update the intent names as described in [Updating a Blueprint for a Specific Locale](#).
- 7. Test the Lambda function.
 - a. Choose **Select a test event, Configure test events**.
 - b. Choose **Amazon Lex Order Flowers** from the **Event template** list. This sample event matches the Amazon Lex request/response model (see [Using Lambda Functions](#)). Give the test event a name (LexOrderFlowersTest).
 - c. Choose **Create**.
 - d. Choose **Test** to test the code hook.
 - e. Verify that the Lambda function ran successfully. The response in this case matches the Amazon Lex response model.

Next Step

[Step 4: Add the Lambda Function as Code Hook \(Console\)](#)

Step 4: Add the Lambda Function as Code Hook (Console)

In this section, you update the configuration of the OrderFlowers intent to use the Lambda function as follows:

- First use the Lambda function as a code hook to perform fulfillment of the OrderFlowers intent. You test the bot and verify that you received a fulfillment message from the Lambda function. Amazon Lex invokes the Lambda function only after you provide data for all the required slots for ordering flowers.
- Configure the same Lambda function as a code hook to perform initialization and validation. You test and verify that the Lambda function performs validation (as you provide slot data).

To add a Lambda function as a code hook (console)

1. In the Amazon Lex console, select the **OrderFlowers** bot. The console shows the **OrderFlowers** intent. Make sure that the intent version is set to \$LATEST because this is the only version that we can modify.

2. Add the Lambda function as the fulfillment code hook and test it.
 - a. In the Editor, choose **AWS Lambda function** as **Fulfillment**, and select the Lambda function that you created in the preceding step (`OrderFlowersCodeHook`). Choose **OK** to give Amazon Lex permission to invoke the Lambda function.

You are configuring this Lambda function as a code hook to fulfill the intent. Amazon Lex invokes this function only after it has all the necessary slot data from the user to fulfill the intent.

- b. Specify a **Goodbye message**.
 - c. Choose **Build**.
 - d. Test the bot using the previous conversation.


The last statement "Thanks, your order for roses....." is a response from the Lambda function that you configured as a code hook. In the preceding section, there was no Lambda function. Now you are using a Lambda function to actually fulfill the `OrderFlowers` intent.

3. Add the Lambda function as an initialization and validation code hook, and test.

The sample Lambda function code that you are using can both perform user input validation and fulfillment. The input event the Lambda function receives has a field (`invocationSource`) that the code uses to determine what portion of the code to run. For more information, see [Lambda Function Input Event and Response Format](#).

- a. Select the `$LATEST` version of the `OrderFlowers` intent. That's is the only version that you can update.
 - b. In the Editor, choose **Initialization and validation in Options**.
 - c. Again, select the same Lambda function.
 - d. Choose **Build**.
 - e. Test the bot.

You are now ready to converse with Amazon Lex as in the following image. To test the validation portion, choose time 6 PM, and your Lambda function returns a response ("Our business hours are from 10 AM to 5 PM."), and prompts you again. After you provide all the valid slot data, the Lambda function fulfills the order.



Next Step

[Step 5 \(Optional\): Review the Details of the Information Flow \(Console\)](#)

Step 5 (Optional): Review the Details of the Information Flow (Console)

This section explains the flow of information between the client and Amazon Lex for each user input, including the integration of the Lambda function.

Note

The section assumes that the client sends requests to Amazon Lex using the `PostText` runtime API and shows request and response details accordingly. For an example of the information flow between the client and Amazon Lex in which client uses the `PostContent` API, see [Step 2a \(Optional\): Review the Details of the Spoken Information Flow \(Console\)](#).

For more information about the `PostText` runtime API and additional details on the requests and responses shown in the following steps, see [PostText](#).

1. User: I would like to order some flowers.
 - a. The client (console) sends the following [PostText](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/ignw84y6seypre4xly5rimopuri2xwnd/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText": "I would like to order some flowers",
  "sessionAttributes": {}
}
```

Both the request URI and the body provide information to Amazon Lex:

- Request URI – Provides bot name (`OrderFlowers`), bot alias (`$LATEST`), and user name (a random string identifying the user). The trailing text indicates that it is a `PostText` API request (and not `PostContent`).
- Request body – Includes the user input (`inputText`) and empty `sessionAttributes`. When the client makes the first request, there are no session attributes. The Lambda function initiates them later.

- b. From the `inputText`, Amazon Lex detects the intent (`OrderFlowers`). This intent is configured with a Lambda function as a code hook for user data initialization and validation. Therefore, Amazon Lex invokes that Lambda function by passing the following information as event data:

```
{
  "messageVersion": "1.0",
  "invocationSource": "DialogCodeHook",
  "userId": "ignw84y6seypre4xly5rimopuri2xwnd",
  "sessionAttributes": {},
  "bot": {
    "name": "OrderFlowers",
    "alias": null,
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "OrderFlowers",
    "slots": {
      "PickupTime": null,
      "FlowerType": null,
      "PickupDate": null
    },
    "confirmationStatus": "None"
  }
}
```

For more information, see [Input Event Format](#).

In addition to the information that the client sent, Amazon Lex also includes the following additional data:

- `messageVersion` – Currently Amazon Lex supports only the 1.0 version.
 - `invocationSource` – Indicates the purpose of Lambda function invocation. In this case, it is to perform user data initialization and validation. At this time, Amazon Lex knows that the user has not provided all the slot data to fulfill the intent.
 - `currentIntent` information with all of the slot values set to null.
- c. At this time, all the slot values are null. There is nothing for the Lambda function to validate. The Lambda function returns the following response to Amazon Lex:

```
{
  "sessionAttributes": {},
  "dialogAction": {
    "type": "Delegate",
    "slots": {
      "PickupTime": null,
      "FlowerType": null,
      "PickupDate": null
    }
  }
}
```

For information about the response format, see [Response Format](#).


Note the following:

- `dialogAction.type` – By setting this value to `Delegate`, Lambda function delegates the responsibility of deciding the next course of action to Amazon Lex.

 Note

If Lambda function detects anything in the user data validation, it instructs Amazon Lex what to do next, as shown in the next few steps.

- d. According to the `dialogAction.type`, Amazon Lex decides the next course of action. Because none of the slots are filled, it decides to elicit the value for the `FlowerType` slot. It selects one of the value elicitation prompts ("What type of flowers would you like to order?") for this slot and sends the following response back to the client:



The client displays the message in the response.

2. User: roses

- a. The client sends the following [PostText](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/ignw84y6seypre4xly5rimopuri2xwnd/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText": "roses",
  "sessionAttributes": {}
}
```

In the request body, the `inputText` provides user input. The `sessionAttributes` remains empty.

- b. Amazon Lex first interprets the `inputText` in the context of the current intent. The service remembers that it had asked the specific user for information about the `FlowerType` slot. It updates the slot value in the current intent and invokes the Lambda function with the following event data:

```
{
  "messageVersion": "1.0",
  "invocationSource": "DialogCodeHook",
}
```

```
"userId": "ignw84y6seypre4xly5rimopuri2xwnd",
"sessionAttributes": {},
"bot": {
  "name": "OrderFlowers",
  "alias": null,
  "version": "$LATEST"
},
"outputDialogMode": "Text",
"currentIntent": {
  "name": "OrderFlowers",
  "slots": {
    "PickupTime": null,
    "FlowerType": "roses",
    "PickupDate": null
  },
  "confirmationStatus": "None"
}
}
```


Note the following:

- `invocationSource` – continues to be `DialogCodeHook` (we are simply validating user data).
 - `currentIntent.slots` – Amazon Lex has updated the `FlowerType` slot to `roses`.
- c. According to the `invocationSource` value of `DialogCodeHook`, the Lambda function performs user data validation. It recognizes `roses` as a valid slot value (and sets `Price` as a session attribute) and returns the following response to Amazon Lex.

```
{
  "sessionAttributes": {
    "Price": 25
  },
  "dialogAction": {
    "type": "Delegate",
    "slots": {
      "PickupTime": null,
      "FlowerType": "roses",
      "PickupDate": null
    }
  }
}
```


Note the following:

- `sessionAttributes` – Lambda function has added `Price` (of the roses) as a session attribute.
 - `dialogAction.type` – is set to `Delegate`. The user data was valid so the Lambda function directs Amazon Lex to choose the next course of action.
- d. According to the `dialogAction.type`, Amazon Lex chooses the next course of action. Amazon Lex knows it needs more slot data so it picks the next unfilled slot (`PickupDate`) with the highest priority according to the intent configuration. Amazon Lex selects one of the value-elicitation prompt messages—"What day do you want the roses to be picked up?"—for this slot according to the intent configuration, and then sends the following response back to the client:



The client simply displays the message in the response – "What day do you want the roses to be picked up?."

3. User: tomorrow

- a. The client sends the following [PostText](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/ignw84y6seypre4xly5rimopuri2xwnd/text
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"
```

```
{
  "inputText": "tomorrow",
  "sessionAttributes": {
    "Price": "25"
  }
}
```

In the request body, `inputText` provides user input and the client passes the session attributes back to the service.

- b. Amazon Lex remembers the context—that it was eliciting data for the `PickupDate` slot. In this context, it knows the `inputText` value is for the `PickupDate` slot. Amazon Lex then invokes the Lambda function by sending the following event:

```
{
  "messageVersion": "1.0",
  "invocationSource": "DialogCodeHook",
  "userId": "ignw84y6seypre4xly5rimopuri2xwnd",
  "sessionAttributes": {
    "Price": "25"
  },
  "bot": {
    "name": "OrderFlowersCustomWithRespCard",
    "alias": null,
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "OrderFlowers",
    "slots": {
      "PickupTime": null,
      "FlowerType": "roses",
      "PickupDate": "2017-01-05"
    },
    "confirmationStatus": "None"
  }
}
```


Amazon Lex has updated the `currentIntent.slots` by setting the `PickupDate` value. Also note that the service passes the `sessionAttributes` as it is to the Lambda function.

- c. As per invocationSource value of DialogCodeHook, the Lambda function performs user data validation. It recognizes PickupDate slot value is valid and returns the following response to Amazon Lex:

```
{
  "sessionAttributes": {
    "Price": 25
  },
  "dialogAction": {
    "type": "Delegate",
    "slots": {
      "PickupTime": null,
      "FlowerType": "roses",
      "PickupDate": "2017-01-05"
    }
  }
}
```

Note the following:

- sessionAttributes – No change.
 - dialogAction.type – is set to Delegate. The user data was valid, and the Lambda function directs Amazon Lex to choose the next course of action.
- d. According to the dialogAction.type, Amazon Lex chooses the next course of action. Amazon Lex knows it needs more slot data so it picks the next unfilled slot (PickupTime) with the highest priority according to the intent configuration. Amazon Lex selects one of the prompt messages ("Deliver the roses at what time on 2017-01-05?") for this slot according to the intent configuration and sends the following response back to the client:



The client displays the message in the response – "Deliver the roses at what time on 2017-01-05?"

4. User: 4 pm

a. The client sends the following [PostText](#) request to Amazon Lex:

```
POST /bot/OrderFlowers/alias/$LATEST/user/ignw84y6seypre4xly5rimopuri2xwnd/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText": "4 pm",
  "sessionAttributes": {
    "Price": "25"
  }
}
```

In the request body, `inputText` provides user input. The client passes the `sessionAttributes` in the request.

b. Amazon Lex understands context. It understands that it was eliciting data for the `PickupTime` slot. In this context, it knows that the `inputText` value is for the `PickupTime` slot. Amazon Lex then invokes the Lambda function by sending the following event:

```
{
  "messageVersion": "1.0",
```

```
"invocationSource": "DialogCodeHook",
"userId": "ignw84y6seypre4xly5rimopuri2xwnd",
"sessionAttributes": {
  "Price": "25"
},
"bot": {
  "name": "OrderFlowersCustomWithRespCard",
  "alias": null,
  "version": "$LATEST"
},
"outputDialogMode": "Text",
"currentIntent": {
  "name": "OrderFlowers",
  "slots": {
    "PickupTime": "16:00",
    "FlowerType": "roses",
    "PickupDate": "2017-01-05"
  },
  "confirmationStatus": "None"
}
}
```


Amazon Lex has updated the `currentIntent.slots` by setting the `PickupTime` value.

- c. According to the `invocationSource` value of `DialogCodeHook`, the Lambda function performs user data validation. It recognizes `PickupDate` slot value is valid and returns the following response to Amazon Lex.

```
{
  "sessionAttributes": {
    "Price": 25
  },
  "dialogAction": {
    "type": "Delegate",
    "slots": {
      "PickupTime": "16:00",
      "FlowerType": "roses",
      "PickupDate": "2017-01-05"
    }
  }
}
```

Note the following:

- `sessionAttributes` – No change in session attribute.
 - `dialogAction.type` – is set to `Delegate`. The user data was valid so the Lambda function directs Amazon Lex to choose the next course of action.
- d. At this time Amazon Lex knows it has all the slot data. This intent is configured with a confirmation prompt. Therefore, Amazon Lex sends the following response to the user asking for confirmation before fulfilling the intent:



The client simply displays the message in the response and waits for the user response.

5. User: Yes

- a. The client sends the following [PostText](#) request to Amazon Lex:

```

POST /bot/OrderFlowers/alias/$LATEST/user/ignw84y6seypre4xly5rimopuri2xwnd/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText": "yes",
  "sessionAttributes": {
    "Price": "25"
  }
}

```

- b. Amazon Lex interprets the `inputText` in the context of confirming the current intent. Amazon Lex understands that the user wants to proceed with the order. This time Amazon Lex invokes the Lambda function to fulfill the intent by sending the following event,

which sets the `invocationSource` to `FulfillmentCodeHook` in the event it sends to the Lambda function. Amazon Lex also sets the `confirmationStatus` to `Confirmed`.

```
{
  "messageVersion": "1.0",
  "invocationSource": "FulfillmentCodeHook",
  "userId": "ignw84y6seypre4xly5rimopuri2xwnd",
  "sessionAttributes": {
    "Price": "25"
  },
  "bot": {
    "name": "OrderFlowersCustomWithRespCard",
    "alias": null,
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "OrderFlowers",
    "slots": {
      "PickupTime": "16:00",
      "FlowerType": "roses",
      "PickupDate": "2017-01-05"
    },
    "confirmationStatus": "Confirmed"
  }
}
```

Note the following:

- `invocationSource` – This time Amazon Lex set this value to `FulfillmentCodeHook`, directing the Lambda function to fulfill the intent.
 - `confirmationStatus` – is set to `Confirmed`.
- c. This time, the Lambda function fulfills the `OrderFlowers` intent, and returns the following response:


```
{
  "sessionAttributes": {
    "Price": "25"
  },
  "dialogAction": {
    "type": "Close",
```

```
    "fulfillmentState": "Fulfilled",
    "message": {
      "contentType": "PlainText",
      "content": "Thanks, your order for roses has been placed and will
be ready for pickup by 16:00 on 2017-01-05"
    }
  }
}
```

Note the following:

- Sets the `dialogAction.type` – The Lambda function sets this value to `Close`, directing Amazon Lex to not expect a user response.
 - `dialogAction.fulfillmentState` – is set to `Fulfilled` and includes an appropriate message to convey to the user.
- d. Amazon Lex reviews the `fulfillmentState` and sends the following response back to the client.

Amazon Lex then returns the following to the client:



Note that:

- `dialogState` – Amazon Lex sets this value to `fulfilled`.
- `message` – is the same message that the Lambda function provided.

The client displays the message.

6. Now test the bot again. To establish a new (user) context, choose the **Clear** link in the test window. Now provide invalid slot data for the `OrderFlowers` intent. This time the Lambda function performs the data validation, resets invalid slot data value to null, and asks Amazon Lex to prompt the user for valid data. For example, try the following:
 - Jasmine as the flower type (it is not one of the supported flower types).
 - Yesterday as the day when you want to pick up the flowers.
 - After placing your order, enter another flower type instead of replying "yes" to confirm the order. In response, the Lambda function updates the `Price` in the session attribute, keeping a running total of flower orders.


The Lambda function also performs the fulfillment activity.

Next Step

[Step 6: Update the Intent Configuration to Add an Utterance \(Console\)](#)

Step 6: Update the Intent Configuration to Add an Utterance (Console)


The `OrderFlowers` bot is configured with only two utterances. This provides limited information for Amazon Lex to build a machine learning model that recognizes and responds to the user's intent. Try typing "I want to order flowers", as in the following test window. Amazon Lex doesn't recognize the text, and responds with "I didn't understand you, what would you like to do?" You can improve the machine learning model by adding more utterances.



Each utterance that you add provides Amazon Lex with more information about how to respond to your users. You don't need to add an exact utterance, Amazon Lex generalizes from the samples that you provide to recognize both exact matches and similar input.


To add an utterance (console)

1. Add the utterance "I want flowers" to the intent by typing it in the **Sample utterances** section of the intent editor, as in the following image, and then clicking the plus icon next to the new utterance.



2. Build your bot to pick up the change. Choose **Build**, and then choose **Build** again.

3. Test your bot to confirm that it recognized the new utterance. In the test window, as in the following image, type "I want to order flowers." Amazon Lex recognizes the phrase and responds with "What type of flowers would you like to order?".



Next Step

[Step 7 \(Optional\): Clean Up \(Console\)](#)

Step 7 (Optional): Clean Up (Console)

Now, delete the resources that you created and clean up your account.

You can delete only resources that are not in use. In general, you should delete resources in the following order:

- Delete bots to free up intent resources.
- Delete intents to free up slot type resources.
- Delete slot types last.

To clean up your account (console)

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. From the list of bots, choose the check box next to **OrderFlowers**.
3. To delete the bot, choose **Delete**, and then choose **Continue** in the confirmation dialog box.
4. In the left pane, choose **Intents**.
5. In the list of intents, choose **OrderFlowersIntent**.
6. To delete the intent, choose **Delete**, and then choose **Continue** in the confirmation dialog box.
7. In the left pane, choose **Slot types**.
8. In the list of slot types, choose **Flowers**.
9. To delete the slot type, choose **Delete**, and then choose **Continue** in the confirmation dialog box.

You have removed all of the Amazon Lex resources that you created and cleaned up your account. If desired, you can use the [Lambda console](#) to delete the Lambda function used in this exercise.

Exercise 2: Create a Custom Amazon Lex Bot

In this exercise, you use the Amazon Lex console to create a custom bot that orders pizza (OrderPizzaBot). You configure the bot by adding a custom intent (OrderPizza), defining custom slot types, and defining the slots required to fulfill a pizza order (pizza crust, size, and so on). For more information about slot types and slots, see [Amazon Lex: How It Works](#).

Topics

- [Step 1: Create a Lambda Function](#)
- [Step 2: Create a Bot](#)
- [Step 3: Build and Test the Bot](#)
- [Step 4 \(Optional\): Clean up](#)

Step 1: Create a Lambda Function

First, create a Lambda function which fulfills a pizza order. You specify this function in your Amazon Lex bot, which you create in the next section.

To create a Lambda function

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. Choose **Create function**.
3. On the **Create function** page, choose **Author from scratch**.

Because you are using custom code provided to you in this exercise to create a Lambda function, you choose author the function from scratch.

Do the following:

- a. Type the name (PizzaOrderProcessor).
 - b. For the **Runtime**, choose the latest version of Node.js.
 - c. For the **Role**, choose **Create new role from template(s)**.
 - d. Enter a new role name (PizzaOrderProcessorRole).
 - e. Choose **Create function**.
4. On the function page, do the following:

In the **Function code** section, choose **Edit code inline**, and then copy the following Node.js function code and paste it in the window.

```
'use strict';

// Close dialog with the customer, reporting fulfillmentState of Failed or
// Fulfilled ("Thanks, your pizza will arrive in 20 minutes")
function close(sessionAttributes, fulfillmentState, message) {
    return {
        sessionAttributes,
        dialogAction: {
            type: 'Close',
            fulfillmentState,
            message,
        },
    };
}

// ----- Events -----
```

```
function dispatch(intentRequest, callback) {
    console.log(`request received for userId=${intentRequest.userId}, intentName=
${intentRequest.currentIntent.name}`);
    const sessionAttributes = intentRequest.sessionAttributes;
    const slots = intentRequest.currentIntent.slots;
    const crust = slots.crust;
    const size = slots.size;
    const pizzaKind = slots.pizzaKind;

    callback(close(sessionAttributes, 'Fulfilled',
        {'contentType': 'PlainText', 'content': `Okay, I have ordered your ${size}
${pizzaKind} pizza on ${crust} crust`}));
}

// ----- Main handler -----

// Route the incoming request based on intent.
// The JSON body of the request is provided in the event slot.
export const handler = (event, context, callback) => {
    try {
        dispatch(event,
            (response) => {
                callback(null, response);
            });
    } catch (err) {
        callback(err);
    }
};
```

5. Choose **Save**.

Test the Lambda Function Using Sample Event Data

In the console, test the Lambda function by using sample event data to manually invoke it.

To test the Lambda function:

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. On the **Lambda function** page, choose the Lambda function (PizzaOrderProcessor).
3. On the function page, in the list of test events, choose **Configure test events**.

4. On the **Configure test event** page, do the following:
 - a. Choose **Create new test event**.
 - b. In the **Event name** field, enter a name for the event (PizzaOrderProcessorTest).
 - c. Copy the following Amazon Lex event into the window.

```
{
  "messageVersion": "1.0",
  "invocationSource": "FulfillmentCodeHook",
  "userId": "user-1",
  "sessionAttributes": {},
  "bot": {
    "name": "PizzaOrderingApp",
    "alias": "$LATEST",
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "OrderPizza",
    "slots": {
      "size": "large",
      "pizzaKind": "meat",
      "crust": "thin"
    },
    "confirmationStatus": "None"
  }
}
```

5. Choose **Create**.

AWS Lambda creates the test and you go back to the function page. Choose **Test** and Lambda runs your Lambda function.

In the result box, choose **Details**. The console displays the following output in the **Execution result** pane.

```
{
  "sessionAttributes": {},
  "dialogAction": {
    "type": "Close",
    "fulfillmentState": "Fulfilled",
```

```
"message": {
  "contentType": "PlainText",
  "content": "Okay, I have ordered your large meat pizza on thin crust."
}
```

Next Step

[Step 2: Create a Bot](#)

Step 2: Create a Bot

In this step, you create a bot to handle pizza orders.

Topics

- [Create the Bot](#)
- [Create an Intent](#)
- [Create Slot Types](#)
- [Configure the Intent](#)
- [Configure the Bot](#)

Create the Bot


Create the `PizzaOrderingBot` bot with the minimum information needed. You add an intent, an action that the user wants to perform, for the bot later.

To create the bot

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Create a bot.
 - a. If you are creating your first bot, choose **Get Started**. Otherwise, choose **Bots**, and then choose **Create**.
 - b. On the **Create your Lex bot** page, choose **Custom bot** and provide the following information:
 - **Bot name:** `PizzaOrderingBot`
 - **Language:** Choose the language and locale for your bot.

- **Output voice:** Salli
 - **Session timeout :** 5 minutes.
 - **COPPA:** Choose the appropriate response.
 - **User utterance storage:** Choose the appropriate response.
- c. Choose **Create**.

The console sends Amazon Lex a request to create a new bot. Amazon Lex sets the bot version to `$LATEST`. After creating the bot, Amazon Lex shows the bot **Editor** tab, as in the following image:



- The bot version, **Latest**, appears next to the bot name in the console. New Amazon Lex resources have `$LATEST` as the version. For more information, see [Versioning and Aliases](#).
- Because you haven't created any intents or slots types, none are listed.
- **Build** and **Publish** are bot-level activities. After you configure the entire bot, you'll learn more about these activities.

Next Step

[Create an Intent](#)

Create an Intent

Now, create the `OrderPizza` intent , an action that the user wants to perform, with the minimum information needed. You add slot types for the intent and then configure the intent later.

To create an intent

1. In the Amazon Lex console, choose the plus sign (+) next to **Intents**, and then choose **Create new intent**.

2. In the **Create intent** dialog box, type the name of the intent (`OrderPizza`), and then choose **Add**.

The console sends a request to Amazon Lex to create the `OrderPizza` intent. In this example you create slots for the intent after you create slot types.

Next Step

[Create Slot Types](#)

Create Slot Types

Create the slot types, or parameter values, that the `OrderPizza` intent uses.

To create slot types

1. In the left menu, choose the plus sign (+) next to **Slot types**.
2. In the **Add slot type** dialog box, add the following:
 - **Slot type name** – Crusts
 - **Description** – Available crusts
 - Choose **Restrict to Slot values and Synonyms**
 - **Value** – Type **thick**. Press tab and in the **Synonym** field type **stuffed**. Choose the plus sign (+). Type **thin** and then choose the plus sign (+) again.

The dialog should look like the following image:

Add slot type

Slot type name

Crusts

Description

Available crusts

Slot Resolution

☐ Expand Values ⓘ

☒ Restrict to Slot values and Synonyms ⓘ

Value ⓘ

e.g. Small

Enter Synonym +

Press Tab to add a synonym

thick

stuffed × | ×

thin

unstuffed ×

Cancel
Save slot type
Add slot to intent

- Choose **Add slot to intent**.
- On the **Intent** page, choose **Required**. Change the name of the slot from **slotOne** to **crust**. Change the prompt to **What kind of crust would you like?**
- Repeat [Step 1](#) through [Step 4](#) using the values in the following table:

Name	Description	Values	Slot name	Prompt
Sizes	Available sizes	small, medium, large	size	What size pizza?
PizzaKind	Available pizzas	veg, cheese	pizzaKind	Do you want a veg or cheese pizza?

Next Step

[Configure the Intent](#)

Configure the Intent

Configure the `OrderPizza` intent to fulfill a user's request to order a pizza.

To configure an intent

- On the **OrderPizza** configuration page, configure the intent as follows:
 - Sample utterances** – Type the following strings. The curly braces `{}` enclose slot names.
 - I want to order pizza please
 - I want to order a pizza
 - I want to order a {pizzaKind} pizza
 - I want to order a {size} {pizzaKind} pizza
 - I want a {size} {crust} crust {pizzaKind} pizza
 - Can I get a pizza please
 - Can I get a {pizzaKind} pizza
 - Can I get a {size} {pizzaKind} pizza
 - Lambda initialization and validation** – Leave the default setting.
 - Confirmation prompt** – Leave the default setting.
 - Fulfillment** – Perform the following tasks:
 - Choose **AWS Lambda function**.
 - Choose **PizzaOrderProcessor**.

- If the **Add permission to Lambda function** dialog box is shown, choose **OK** to give the OrderPizza intent permission to call the PizzaOrderProcessor Lambda function.
- Leave **None** selected.

The intent should look like the following:

OrderPizza Latest ▾

▼ Sample utterances ⓘ

e.g. I would like to book a flight. +

I want to order a pizza please ✕

I want to order a pizza ✕

I want to order a {pizzaKind} pizza ✕

I want to order a {size} {pizzaKind} pizza ✕

I want to order a {size} {crust} crust {pizzaKind} pizza ✕

Can I get a pizza please ✕

Can I get a {pizzaKind} pizza ✕

Can I get a {size} {pizzaKind} pizza ✕

► Lambda initialization and validation ⓘ

▼ Slots ⓘ

Priority	Required	Name	Slot type	Prompt		
		<i>e.g. Location</i>	<i>e.g. AMAZO...</i> ▾	<i>e.g. What city?</i>	⚙	+
1. ▾	☒	crust	Crusts ▾	1 ▾	What kind of crust would you	✕
2. ^ ▾	☒	size	Sizes ▾	1 ▾	What size pizza	✕
3. ^	☒	pizzaKind	PizzaKind ▾	1 ▾	Do you want a veg or chees	✕

► Confirmation prompt ⓘ

▼ Fulfillment ⓘ

☒ AWS Lambda function ☐ Return parameters to client

PizzaOrderProcessor ▾


Next Step

[Configure the Bot](#)

Configure the Bot

Configure error handling for the PizzaOrderingBot bot.

1. Navigate to the PizzaOrderingBot bot. Choose **Editor**, and then choose **Error Handling**, as in the following image:



2. Use the **Editor** tab to configure bot error handling.

- Information you provide in **Clarification Prompts** maps to the bot's [clarificationPrompt](#) configuration.

When Amazon Lex can't determine the user intent, the service returns a response with this message.

- Information that you provide in the **Hang-up** phrase maps to the bot's [abortStatement](#) configuration.

If the service can't determine the user's intent after a set number of consecutive requests, Amazon Lex returns a response with this message.

Leave the defaults.

Next Step

[Step 3: Build and Test the Bot](#)

Step 3: Build and Test the Bot

Make sure the bot works, by building and testing it.


To build and test the bot

1. To build the PizzaOrderingBot bot, choose **Build**.

Amazon Lex builds a machine learning model for the bot. When you test the bot, the console uses the runtime API to send the user input back to Amazon Lex. Amazon Lex then uses the machine learning model to interpret the user input.

It can take some time to complete the build.

2. To test the bot, in the **Test Bot** window, start communicating with your Amazon Lex bot.
 - For example, you might say or type the following:



- Use the sample utterances that you configured in the `OrderPizza` intent to test the bot. For example, the following is one of the sample utterances that you configured for the `PizzaOrder` intent:

I want a {size} {crust} crust {pizzaKind} pizza

To test it, type the following:

I want a large thin crust cheese pizza

When you type "I want to order a pizza," Amazon Lex detects the intent (`OrderPizza`). Then, Amazon Lex asks for slot information.


After you provide all of the slot information, Amazon Lex invokes the Lambda function that you configured for the intent.

The Lambda function returns a message ("Okay, I have ordered your ...") to Amazon Lex, which Amazon Lex returns to you..

Inspecting the Response

Underneath the chat window is a pane that enables you to inspect the response from Amazon Lex. The pane provides comprehensive information about the state of your bot that changes as you interact with your bot. The contents of the panes show you the current state of the operation.

- **Dialog State** – The current state of the conversation with the user. It can be `ElicitIntent`, `ElicitSlot`, `ConfirmIntent` or `Fulfilled`.
- **Summary** – Shows a simplified view of the dialog that shows the slot values for the intent being fulfilled so that you can keep track of the information flow. It shows the intent name, the number of slots and the number of slots filled, and a list of all of the slots and their associated values. See the following image:



- **Detail** – Shows the raw JSON response from the chatbot to give you a deeper view into the bot interaction and the current state of the dialog as you test and debug your chatbot. If you type in the chat window, the inspection pane shows the JSON response from the [PostText](#) operation. If you speak to the chat window, the inspection pane shows the response headers from the [PostContent](#) operation. See the following image:

Inspect Response

Dialog State: ElicitSlot

☐ Summary ☒ Detail

RequestID: 41392c21-97ff-11e7-a10b-5bcc0093a006

```
{
  "dialogState": "ElicitsSlot",
  "intentName": "OrderPizza",
  "message": "What kind of crust would you like",
  "responseCard": null,
  "sessionAttributes": {},
  "slotToElicit": "crust",
  "slots": {
    "crust": null,
    "pizzaKind": "cheese",
    "size": "large"
  }
}
```

Next Step

[Step 4 \(Optional\): Clean up](#)

Step 4 (Optional): Clean up

Delete the resources that you created and clean up your account to avoid incurring more charges for the resources you created.

You can delete only resources that are not in use. For example, you cannot delete a slot type that is referenced by an intent. You cannot delete an intent that is referenced by a bot.

Delete resources in the following order:

- Delete bots to free up intent resources.

- Delete intents to free up slot type resources.
- Delete slot types last.

To clean up your account

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. From the list of bots, choose **PizzaOrderingBot**.
3. To delete the bot, choose **Delete**, and then choose **Continue**.
4. In the left pane, choose **Intents**.
5. In the list of intents, choose **OrderPizza**.
6. To delete the intent, choose **Delete**, and then choose **Continue**.
7. In the left menu, choose **Slot types**.
8. In the list of slot types, choose **Crusts**.
9. To delete the slot type, choose **Delete**, and then choose **Continue**.
10. Repeat [Step 8](#) and [Step 9](#) for the Sizes and PizzaKind slot types.

You have removed all of the resources that you created and cleaned up your account.

Next Steps

- [Publish a Version and Create an Alias](#)
- [Create an Amazon Lex bot with the AWS Command Line Interface](#)

Exercise 3: Publish a Version and Create an Alias

In Getting Started Exercises 1 and 2, you created a bot and tested it. In this exercise, you do the following:

- Publish a new version of the bot. Amazon Lex takes a snapshot copy of the \$LATEST version to publish a new version.
- Create an alias that points to the new version.


For more information about versioning and aliases, see [Versioning and Aliases](#).

Do the following to publish a version of a bot you created for this exercise:

1. In the Amazon Lex console, choose one of the bots you created.

Verify that the console shows the `$LATEST` as the bot version next to the bot name.

2. Choose **Publish**.
3. On the **Publish *botname*** wizard, specify the alias **BETA**, and then choose **Publish**.
4. Verify that the Amazon Lex console shows the new version next to the bot name, as in the following image.



Now that you have a working bot with published version and an alias, you can deploy the bot (in your mobile application or integrate the bot with Facebook Messenger). For an example, see [Integrating an Amazon Lex Bot with Facebook Messenger](#).

Step 4: Getting Started (AWS CLI)

In this step, you use the AWS CLI to create, test, and modify an Amazon Lex bot. To complete these exercises, you need to be familiar with using the CLI and have a text editor. For more information, see [Step 2: Set Up the AWS Command Line Interface](#)

- Exercise 1 — Create and test an Amazon Lex bot. The exercise provides all of the JSON objects that you need to create a custom slot type, an intent, and a bot. For more information, see [Amazon Lex: How It Works](#)

- Exercise 2 — Update the bot that you created in Exercise 1 to add an additional sample utterance. Amazon Lex uses sample utterances to build the machine learning model for your bot.
- Exercise 3 — Update the bot that you created in Exercise 1 to add a Lambda function to validate user input and to fulfill the intent.
- Exercise 4 — Publish a version of the slot type, intent, and bot resources that you created in Exercise 1. A version is a snapshot of a resource that can't be changed.
- Exercise 5 — Create an alias for the bot that you created in Exercise 1.
- Exercise 6 — Clean up your account by deleting the slot type, intent, and bot that you created in Exercise 1, and the alias that you created in Exercise 5.

Topics

- [Exercise 1: Create an Amazon Lex Bot \(AWS CLI\)](#)
- [Exercise 2: Add a New Utterance \(AWS CLI\)](#)
- [Exercise 3: Add a Lambda Function \(AWS CLI\)](#)
- [Exercise 4: Publish a Version \(AWS CLI\)](#)
- [Exercise 5: Create an Alias \(AWS CLI\)](#)
- [Exercise 6: Clean Up \(AWS CLI\)](#)

Exercise 1: Create an Amazon Lex Bot (AWS CLI)

In general, when you create bots, you:

1. Create slot types to define the information that your bot will be working with.
2. Create intents that define the user actions that your bot supports. Use the custom slot types that you created earlier to define the slots, or parameters, that your intent requires.
3. Create a bot that uses the intents that you defined.

In this exercise you create and test a new Amazon Lex bot using the CLI. Use the JSON structures that we provide to create the bot. To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#).

Topics

- [Step 1: Create a Service-Linked Role \(AWS CLI\)](#)

- [Step 2: Create a Custom Slot Type \(AWS CLI\)](#)
- [Step 3: Create an Intent \(AWS CLI\)](#)
- [Step 4: Create a Bot \(AWS CLI\)](#)
- [Step 5: Test a Bot \(AWS CLI\)](#)

Step 1: Create a Service-Linked Role (AWS CLI)

Amazon Lex assumes AWS Identity and Access Management service-linked roles to call AWS services on behalf of your bots. The roles, which are in your account, are linked to Amazon Lex use cases and have predefined permissions. For more information, see [Using Service-Linked Roles for Amazon Lex](#).

If you've already created an Amazon Lex bot using the console, the service-linked role was created automatically. Skip to [Step 2: Create a Custom Slot Type \(AWS CLI\)](#).

To create a service-linked role (AWS CLI)

1. In the AWS CLI, type the following command:

```
aws iam create-service-linked-role --aws-service-name lex.amazonaws.com
```

2. Check the policy using the following command:

```
aws iam get-role --role-name AWSServiceRoleForLexBots
```

The response is:

```
{
  "Role": {
    "AssumeRolePolicyDocument": {
      "Version": "2012-10-17",
      "Statement": [
        {
          "Action": "sts:AssumeRole",
          "Effect": "Allow",
          "Principal": {
            "Service": "lex.amazonaws.com"
          }
        }
      ]
    }
  }
}
```

```
    },
    "RoleName": "AWSServiceRoleForLexBots",
    "Path": "/aws-service-role/lex.amazonaws.com/",
    "Arn": "arn:aws:iam::account-id:role/aws-service-role/lex.amazonaws.com/
AWSServiceRoleForLexBots"
}
```

Next Step

[Step 2: Create a Custom Slot Type \(AWS CLI\)](#)

Step 2: Create a Custom Slot Type (AWS CLI)

Create a custom slot type with enumeration values for the flowers that can be ordered. You use this type in the next step when you create the `OrderFlowers` intent. A *slot type* defines the possible values for a slot, or parameter, of the intent.

To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#).

To create a custom slot type (AWS CLI)

1. Create a text file named **FlowerTypes.json**. Copy the JSON code from [FlowerTypes.json](#) into the text file.
2. Call the [PutSlotType](#) operation using the AWS CLI to create the slot type. The example is formatted for Unix, Linux, and macOS. For Windows, replace the backslash (\) Unix continuation character at the end of each line with a caret (^).

```
aws lex-models put-slot-type \
  --region region \
  --name FlowerTypes \
  --cli-input-json file://FlowerTypes.json
```

The response from the server is:

```
{
  "enumerationValues": [
    {
      "value": "tulips"
    },
  ],
}
```



```
{
  {
    "value": "lilies"
  },
  {
    "value": "roses"
  }
],
"name": "FlowerTypes",
"checksum": "checksum",
"version": "$LATEST",
"lastUpdatedDate": timestamp,
"createdDate": timestamp,
"description": "Types of flowers to pick up"
}
```

Next Step

[Step 3: Create an Intent \(AWS CLI\)](#)

FlowerTypes.json

The following code is the JSON data required to create the FlowerTypes custom slot type:

```
{
  "enumerationValues": [
    {
      "value": "tulips"
    },
    {
      "value": "lilies"
    },
    {
      "value": "roses"
    }
  ],
  "name": "FlowerTypes",
  "description": "Types of flowers to pick up"
}
```

Step 3: Create an Intent (AWS CLI)

Create an intent for the OrderFlowersBot bot and provide three slots, or parameters. The slots allow the bot to fulfill the intent:

- FlowerType is a custom slot type that specifies which types of flowers can be ordered.
- AMAZON.DATE and AMAZON.TIME are built-in slot types used for getting the date and time to deliver the flowers from the user.

To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#).

To create the OrderFlowers intent (AWS CLI)

1. Create a text file named **OrderFlowers.json**. Copy the JSON code from [OrderFlowers.json](#) into the text file.
2. In the AWS CLI, call the [PutIntent](#) operation to create the intent. The example is formatted for Unix, Linux, and macOS. For Windows, replace the backslash (\) Unix continuation character at the end of each line with a caret (^).

```
aws lex-models put-intent \  
  --region region \  
  --name OrderFlowers \  
  --cli-input-json file://OrderFlowers.json
```

The server responds with the following:

```
{  
  "confirmationPrompt": {  
    "maxAttempts": 2,  
    "messages": [  
      {  
        "content": "Okay, your {FlowerType} will be ready for pickup by  
{PickupTime} on {PickupDate}. Does this sound okay?",  
        "contentType": "PlainText"  
      }  
    ]  
  },  
  "name": "OrderFlowers",  
  "checksum": "checksum",
```

```

"version": "$LATEST",
"rejectionStatement": {
  "messages": [
    {
      "content": "Okay, I will not place your order.",
      "contentType": "PlainText"
    }
  ]
},
"createdDate": timestamp,
"lastUpdatedDate": timestamp,
"sampleUtterances": [
  "I would like to pick up flowers",
  "I would like to order some flowers"
],
"slots": [
  {
    "slotType": "AMAZON.TIME",
    "name": "PickupTime",
    "slotConstraint": "Required",
    "valueElicitationPrompt": {
      "maxAttempts": 2,
      "messages": [
        {
          "content": "Pick up the {FlowerType} at what time on
{PickupDate}?",
          "contentType": "PlainText"
        }
      ]
    },
    "priority": 3,
    "description": "The time to pick up the flowers"
  },
  {
    "slotType": "FlowerTypes",
    "name": "FlowerType",
    "slotConstraint": "Required",
    "valueElicitationPrompt": {
      "maxAttempts": 2,
      "messages": [
        {
          "content": "What type of flowers would you like to
order?",
          "contentType": "PlainText"
        }
      ]
    }
  }
]

```

```

        }
      ]
    },
    "priority": 1,
    "slotTypeVersion": "$LATEST",
    "sampleUtterances": [
      "I would like to order {FlowerType}"
    ],
    "description": "The type of flowers to pick up"
  },
  {
    "slotType": "AMAZON.DATE",
    "name": "PickupDate",
    "slotConstraint": "Required",
    "valueElicitationPrompt": {
      "maxAttempts": 2,
      "messages": [
        {
          "content": "What day do you want the {FlowerType} to be
picked up?",
          "contentType": "PlainText"
        }
      ]
    },
    "priority": 2,
    "description": "The date to pick up the flowers"
  }
],
"fulfillmentActivity": {
  "type": "ReturnIntent"
},
"description": "Intent to order a bouquet of flowers for pick up"
}

```

Next Step

[Step 4: Create a Bot \(AWS CLI\)](#)

OrderFlowers.json

The following code is the JSON data required to create the OrderFlowers intent:

```
{
```

```
"confirmationPrompt": {
  "maxAttempts": 2,
  "messages": [
    {
      "content": "Okay, your {FlowerType} will be ready for pickup by
{PickupTime} on {PickupDate}. Does this sound okay?",
      "contentType": "PlainText"
    }
  ]
},
"name": "OrderFlowers",
"rejectionStatement": {
  "messages": [
    {
      "content": "Okay, I will not place your order.",
      "contentType": "PlainText"
    }
  ]
},
"sampleUtterances": [
  "I would like to pick up flowers",
  "I would like to order some flowers"
],
"slots": [
  {
    "slotType": "FlowerTypes",
    "name": "FlowerType",
    "slotConstraint": "Required",
    "valueElicitationPrompt": {
      "maxAttempts": 2,
      "messages": [
        {
          "content": "What type of flowers would you like to order?",
          "contentType": "PlainText"
        }
      ]
    },
    "priority": 1,
    "slotTypeVersion": "$LATEST",
    "sampleUtterances": [
      "I would like to order {FlowerType}"
    ],
    "description": "The type of flowers to pick up"
  },

```

```

    {
      "slotType": "AMAZON.DATE",
      "name": "PickupDate",
      "slotConstraint": "Required",
      "valueElicitationPrompt": {
        "maxAttempts": 2,
        "messages": [
          {
            "content": "What day do you want the {FlowerType} to be picked
up?",
            "contentType": "PlainText"
          }
        ]
      },
      "priority": 2,
      "description": "The date to pick up the flowers"
    },
    {
      "slotType": "AMAZON.TIME",
      "name": "PickupTime",
      "slotConstraint": "Required",
      "valueElicitationPrompt": {
        "maxAttempts": 2,
        "messages": [
          {
            "content": "Pick up the {FlowerType} at what time on
{PickupDate}?",
            "contentType": "PlainText"
          }
        ]
      },
      "priority": 3,
      "description": "The time to pick up the flowers"
    }
  ],
  "fulfillmentActivity": {
    "type": "ReturnIntent"
  },
  "description": "Intent to order a bouquet of flowers for pick up"
}

```

Step 4: Create a Bot (AWS CLI)

The OrderFlowersBot bot has one intent, the OrderFlowers intent that you created in the previous step. To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#).

Note

The following AWS CLI example is formatted for Unix, Linux, and macOS. For Windows, change "\$LATEST" to \$LATEST.

To create the OrderFlowersBot bot (AWS CLI)

1. Create a text file named **OrderFlowersBot.json**. Copy the JSON code from [OrderFlowersBot.json](#) into the text file.
2. In the AWS CLI, call the [PutBot](#) operation to create the bot. The example is formatted for Unix, Linux, and macOS. For Windows, replace the backslash (\) Unix continuation character at the end of each line with a caret (^).

```
aws lex-models put-bot \  
  --region region \  
  --name OrderFlowersBot \  
  --cli-input-json file://OrderFlowersBot.json
```

The response from the server follows. When you create or update bot, the status field is set to BUILDING. This indicates that the bot isn't ready to use. To determine when the bot is ready for use, use the [GetBot](#) operation in the next step.

```
{  
  "status": "BUILDING",  
  "intents": [  
    {  
      "intentVersion": "$LATEST",  
      "intentName": "OrderFlowers"  
    }  
  ],  
  "name": "OrderFlowersBot",  
  "locale": "en-US",  
  "checksum": "checksum",
```

```

    "abortStatement": {
      "messages": [
        {
          "content": "Sorry, I'm not able to assist at this time",
          "contentType": "PlainText"
        }
      ]
    },
    "version": "$LATEST",
    "lastUpdatedDate": timestamp,
    "createdDate": timestamp,
    "clarificationPrompt": {
      "maxAttempts": 2,
      "messages": [
        {
          "content": "I didn't understand you, what would you like to do?",
          "contentType": "PlainText"
        }
      ]
    },
    "voiceId": "Salli",
    "childDirected": false,
    "idleSessionTTLInSeconds": 600,
    "processBehavior": "BUILD",
    "description": "Bot to order flowers on the behalf of a user"
  }
}

```

3. To determine if your new bot is ready for use, run the following command. Repeat this command until the status field returns READY. The example is formatted for Unix, Linux, and macOS. For Windows, replace the backslash (\) Unix continuation character at the end of each line with a caret (^).

```

aws lex-models get-bot \
  --region region \
  --name OrderFlowersBot \
  --version-or-alias "$LATEST"

```

Look for the status field in the response:

```

{
  "status": "READY",

```



```
...  
}
```

Next Step

[Step 5: Test a Bot \(AWS CLI\)](#)

OrderFlowersBot.json

The following code provides the JSON data required to build the OrderFlowers Amazon Lex bot:

```
{  
  "intents": [  
    {  
      "intentVersion": "$LATEST",  
      "intentName": "OrderFlowers"  
    }  
  ],  
  "name": "OrderFlowersBot",  
  "locale": "en-US",  
  "abortStatement": {  
    "messages": [  
      {  
        "content": "Sorry, I'm not able to assist at this time",  
        "contentType": "PlainText"  
      }  
    ]  
  },  
  "clarificationPrompt": {  
    "maxAttempts": 2,  
    "messages": [  
      {  
        "content": "I didn't understand you, what would you like to do?",  
        "contentType": "PlainText"  
      }  
    ]  
  },  
  "voiceId": "Salli",  
  "childDirected": false,  
  "idleSessionTTLInSeconds": 600,  
  "description": "Bot to order flowers on the behalf of a user"
```

```
}
```

Step 5: Test a Bot (AWS CLI)

To test the bot, you can use either a text-based or a speech-based test.

Topics

- [Test the Bot Using Text Input \(AWS CLI\)](#)
- [Test the Bot Using Speech Input \(AWS CLI\)](#)

Test the Bot Using Text Input (AWS CLI)

To verify that the bot works correctly with text input, use the [PostText](#) operation. To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Runtime Service Quotas](#).

Note

The following AWS CLI example is formatted for Unix, Linux, and macOS. For Windows, change "\$LATEST" to \$LATEST and replace the backslash (\) continuation character at the end of each line with a caret (^).

To use text to test the bot (AWS CLI)

1. In the AWS CLI, start a conversation with the OrderFlowersBot bot. The example is formatted for Unix, Linux, and macOS. For Windows, replace the backslash (\) Unix continuation character at the end of each line with a caret (^).

```
aws lex-runtime post-text \  
  --region region \  
  --bot-name OrderFlowersBot \  
  --bot-alias "$LATEST" \  
  --user-id UserOne \  
  --input-text "i would like to order flowers"
```

Amazon Lex recognizes the user's intent and starts a conversation by returning the following response:

```
{
  "slotToElicit": "FlowerType",
  "slots": {
    "PickupDate": null,
    "PickupTime": null,
    "FlowerType": null
  },
  "dialogState": "ElicitSlot",
  "message": "What type of flowers would you like to order?",
  "intentName": "OrderFlowers"
}
```

2. Run the following commands to finish the conversation with the bot.

```
aws lex-runtime post-text \
  --region region \
  --bot-name OrderFlowersBot \
  --bot-alias "\$LATEST" \
  --user-id UserOne \
  --input-text "roses"
```

```
aws lex-runtime post-text \
  --region region \
  --bot-name OrderFlowersBot \
  --bot-alias "\$LATEST" \
  --user-id UserOne \
  --input-text "tuesday"
```

```
aws lex-runtime post-text \
  --region region \
  --bot-name OrderFlowersBot --bot-alias "\$LATEST" \
  --user-id UserOne \
  --input-text "10:00 a.m."
```

```
aws lex-runtime post-text \
  --region region \
  --bot-name OrderFlowersBot \
  --bot-alias "\$LATEST" \
  --user-id UserOne \
  --input-text "yes"
```

After you confirm the order, Amazon Lex sends a fulfillment response to complete the conversation:

```
{
  "slots": {
    "PickupDate": "2017-05-16",
    "PickupTime": "10:00",
    "FlowerType": "roses"
  },
  "dialogState": "ReadyForFulfillment",
  "intentName": "OrderFlowers"
}
```

Next Step

[Test the Bot Using Speech Input \(AWS CLI\)](#)

Test the Bot Using Speech Input (AWS CLI)

To test the bot using audio files, use the [PostContent](#) operation. You generate the audio files using Amazon Polly text-to-speech operations.

To run the commands in this exercise, you need to know the region the Amazon Lex and Amazon Polly commands will be run. For a list of regions for Amazon Lex, see [Runtime Service Quotas](#). For a list of regions for Amazon Polly see [AWS Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

Note

The following AWS CLI example is formatted for Unix, Linux, and macOS. For Windows, change "\$LATEST" to \$LATEST and replace the backslash (\) continuation character at the end of each line with a caret (^).

To use a speech input to test the bot (AWS CLI)

1. In the AWS CLI, create an audio file using Amazon Polly. The example is formatted for Unix, Linux, and macOS. For Windows, replace the backslash (\) Unix continuation character at the end of each line with a caret (^).

```
aws polly synthesize-speech \  
  --region region \  
  --output-format pcm \  
  --text "i would like to order flowers" \  
  --voice-id "Salli" \  
  IntentSpeech.mpg
```

2. To send the audio file to Amazon Lex, run the following command. Amazon Lex saves the audio from the response in the specified output file.

```
aws lex-runtime post-content \  
  --region region \  
  --bot-name OrderFlowersBot \  
  --bot-alias "\$LATEST" \  
  --user-id UserOne \  
  --content-type "audio/l16; rate=16000; channels=1" \  
  --input-stream IntentSpeech.mpg \  
  IntentOutputSpeech.mpg
```

Amazon Lex responds with a request for the first slot. It saves the audio response in the specified output file.

```
{  
  "contentType": "audio/mpeg",  
  "slotToElicit": "FlowerType",  
  "dialogState": "ElicitSlot",  
  "intentName": "OrderFlowers",  
  "inputTranscript": "i would like to order some flowers",  
  "slots": {  
    "PickupDate": null,  
    "PickupTime": null,  
    "FlowerType": null  
  },  
  "message": "What type of flowers would you like to order?"  
}
```

3. To order roses, create the following audio file and send it to Amazon Lex :

```
aws polly synthesize-speech \  
  --region region \  
  --output-format pcm \  
  IntentSpeech.mpg
```

```
--text "roses" \  
--voice-id "Salli" \  
FlowerTypeSpeech.mpg
```

```
aws lex-runtime post-content \  
  --region region \  
  --bot-name OrderFlowersBot \  
  --bot-alias "\$LATEST" \  
  --user-id UserOne \  
  --content-type "audio/l16; rate=16000; channels=1" \  
  --input-stream FlowerTypeSpeech.mpg \  
  FlowerTypeOutputSpeech.mpg
```

4. To set the delivery date, create the following audio file and send it to Amazon Lex:

```
aws polly synthesize-speech \  
  --region region \  
  --output-format pcm \  
  --text "tuesday" \  
  --voice-id "Salli" \  
  DateSpeech.mpg
```

```
aws lex-runtime post-content \  
  --region region \  
  --bot-name OrderFlowersBot \  
  --bot-alias "\$LATEST" \  
  --user-id UserOne \  
  --content-type "audio/l16; rate=16000; channels=1" \  
  --input-stream DateSpeech.mpg \  
  DateOutputSpeech.mpg
```

5. To set the delivery time, create the following audio file and send it to Amazon Lex:

```
aws polly synthesize-speech \  
  --region region \  
  --output-format pcm \  
  --text "10:00 a.m." \  
  --voice-id "Salli" \  
  TimeSpeech.mpg
```

```
aws lex-runtime post-content \  
  --region region \  
  --bot-name OrderFlowersBot \  
  --bot-alias "\$LATEST" \  
  --user-id UserOne \  
  --content-type "audio/l16; rate=16000; channels=1" \  
  --input-stream TimeSpeech.mpg \  
  TimeOutputSpeech.mpg
```

```
--region region \  
--bot-name OrderFlowersBot \  
--bot-alias "$LATEST" \  
--user-id UserOne \  
--content-type "audio/l16; rate=16000; channels=1" \  
--input-stream TimeSpeech.mpg \  
TimeOutputSpeech.mpg
```

6. To confirm the delivery, create the following audio file and send it to Amazon Lex:

```
aws polly synthesize-speech \  
--region region \  
--output-format pcm \  
--text "yes" \  
--voice-id "Salli" \  
ConfirmSpeech.mpg
```

```
aws lex-runtime post-content \  
--region region \  
--bot-name OrderFlowersBot \  
--bot-alias "$LATEST" \  
--user-id UserOne \  
--content-type "audio/l16; rate=16000; channels=1" \  
--input-stream ConfirmSpeech.mpg \  
ConfirmOutputSpeech.mpg
```

After you confirm the delivery, Amazon Lex sends a response that confirms fulfillment of the intent:

```
{  
  "contentType": "text/plain;charset=utf-8",  
  "dialogState": "ReadyForFulfillment",  
  "intentName": "OrderFlowers",  
  "inputTranscript": "yes",  
  "slots": {  
    "PickupDate": "2017-05-16",  
    "PickupTime": "10:00",  
    "FlowerType": "roses"  
  }  
}
```

Next Step

[Exercise 2: Add a New Utterance \(AWS CLI\)](#)

Exercise 2: Add a New Utterance (AWS CLI)

To improve the machine learning model that Amazon Lex uses to recognize requests from your users, add another sample utterance to the bot.

Adding a new utterance is a four-step process.

1. Use the [GetIntent](#) operation to get an intent from Amazon Lex.
2. Update the intent.
3. Use the [PutIntent](#) operation to send the updated intent back to Amazon Lex.
4. Use the [GetBot](#) and [PutBot](#) operations to rebuild any bot that uses the intent.

To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#).

The response from the `GetIntent` operation contains a field called `checksum` that identifies a specific revision of the intent. You must provide the checksum value when you use the [PutIntent](#) operation to update an intent. If you don't, you'll get the following error message:

```
An error occurred (PreconditionFailedException) when calling
the PutIntent operation: Intent intent name already exists.
If you are trying to update intent name you must specify the
checksum.
```

Note

The following AWS CLI example is formatted for Unix, Linux, and macOS. For Windows, change "`\$LATEST`" to `$LATEST` and replace the backslash (`\`) continuation character at the end of each line with a caret (`^`).

To update the OrderFlowers intent (AWS CLI)

1. In the AWS CLI, get the intent from Amazon Lex. Amazon Lex sends the output to a file called **OrderFlowers-V2.json**.

```
aws lex-models get-intent \  
  --region region \  
  --name OrderFlowers \  
  --intent-version "$LATEST" > OrderFlowers-V2.json
```

2. Open **OrderFlowers-V2.json** in a text editor.

1. Find and delete the `createdAt`, `lastUpdatedDate`, and `version` fields.
2. Add the following to the `sampleUtterances` field:

```
I want to order flowers
```

3. Save the file.
3. Send the updated intent to Amazon Lex with the following command:

```
aws lex-models put-intent \  
  --region region \  
  --name OrderFlowers \  
  --cli-input-json file://OrderFlowers-V2.json
```

Amazon Lex sends the following response:

```
{  
  "confirmationPrompt": {  
    "maxAttempts": 2,  
    "messages": [  
      {  
        "content": "Okay, your {FlowerType} will be ready for pickup by  
{PickupTime} on {PickupDate}. Does this sound okay?",  
        "contentType": "PlainText"  
      }  
    ]  
  },  
  "name": "OrderFlowers",  
  "checksum": "checksum",  
  "version": "$LATEST",
```

```

    "rejectionStatement": {
      "messages": [
        {
          "content": "Okay, I will not place your order.",
          "contentType": "PlainText"
        }
      ]
    },
    "createdDate": timestamp,
    "lastUpdatedDate": timestamp,
    "sampleUtterances": [
      "I would like to pick up flowers",
      "I would like to order some flowers",
      "I want to order flowers"
    ],
    "slots": [
      {
        "slotType": "AMAZON.TIME",
        "name": "PickupTime",
        "slotConstraint": "Required",
        "valueElicitationPrompt": {
          "maxAttempts": 2,
          "messages": [
            {
              "content": "Pick up the {FlowerType} at what time on
{PickupDate}?",
              "contentType": "PlainText"
            }
          ]
        },
        "priority": 3,
        "description": "The time to pick up the flowers"
      },
      {
        "slotType": "FlowerTypes",
        "name": "FlowerType",
        "slotConstraint": "Required",
        "valueElicitationPrompt": {
          "maxAttempts": 2,
          "messages": [
            {
              "content": "What type of flowers would you like to
order?",
              "contentType": "PlainText"
            }
          ]
        }
      }
    ]
  }
}

```

```

        }
    ]
},
"priority": 1,
"slotTypeVersion": "$LATEST",
"sampleUtterances": [
    "I would like to order {FlowerType}"
],
"description": "The type of flowers to pick up"
},
{
    "slotType": "AMAZON.DATE",
    "name": "PickupDate",
    "slotConstraint": "Required",
    "valueElicitationPrompt": {
        "maxAttempts": 2,
        "messages": [
            {
                "content": "What day do you want the {FlowerType} to be
picked up?",
                "contentType": "PlainText"
            }
        ]
    },
    "priority": 2,
    "description": "The date to pick up the flowers"
}
],
"fulfillmentActivity": {
    "type": "ReturnIntent"
},
"description": "Intent to order a bouquet of flowers for pick up"
}

```

Now that you have updated the intent, rebuild any bot that uses it.

To rebuild the OrderFlowersBot bot (AWS CLI)

1. In the AWS CLI, get the definition of the OrderFlowersBot bot and save it to a file with the following command:

```
aws lex-models get-bot \
```

```
--region region \  
--name OrderFlowersBot \  
--version-or-alias "$LATEST" > OrderFlowersBot-V2.json
```

2. In a text editor, open **OrderFlowersBot-V2.json**. Remove the `createdDate`, `lastUpdatedDate`, `status` and `version` fields.
3. In a text editor, add the following line to the bot definition:

```
"processBehavior": "BUILD",
```

4. In the AWS CLI, build a new revision of the bot by running the following command to :

```
aws lex-models put-bot \  
--region region \  
--name OrderFlowersBot \  
--cli-input-json file://OrderFlowersBot-V2.json
```

The response from the server is:

```
{  
  "status": "BUILDING",  
  "intents": [  
    {  
      "intentVersion": "$LATEST",  
      "intentName": "OrderFlowers"  
    }  
  ],  
  "name": "OrderFlowersBot",  
  "locale": "en-US",  
  "checksum": "checksum",  
  "abortStatement": {  
    "messages": [  
      {  
        "content": "Sorry, I'm not able to assist at this time",  
        "contentType": "PlainText"  
      }  
    ]  
  },  
  "version": "$LATEST",  
  "lastUpdatedDate": timestamp,  
  "createdDate": timestamp  
  "clarificationPrompt": {
```

```
    "maxAttempts": 2,
    "messages": [
      {
        "content": "I didn't understand you, what would you like to do?",
        "contentType": "PlainText"
      }
    ],
    "voiceId": "Salli",
    "childDirected": false,
    "idleSessionTTLInSeconds": 600,
    "description": "Bot to order flowers on the behalf of a user"
  }
}
```

Next Step

[Exercise 3: Add a Lambda Function \(AWS CLI\)](#)

Exercise 3: Add a Lambda Function (AWS CLI)

Add a Lambda function that validates user input and fulfills the user's intent to the bot.

Adding a Lambda expression is a five-step process.

1. Use the Lambda [AddPermission](#) function to enable the OrderFlowers intent to call the Lambda [Invoke](#) operation.
2. Use the [GetIntent](#) operation to get the intent from Amazon Lex.
3. Update the intent to add the Lambda function.
4. Use the [PutIntent](#) operation to send the updated intent back to Amazon Lex.
5. Use the [GetBot](#) and [PutBot](#) operations to rebuild any bot that uses the intent.

To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#).

If you add a Lambda function to an intent before you add the InvokeFunction permission, you get the following error message:

An error occurred (BadRequestException) when calling the

PutIntent operation: Lex is unable to access the Lambda function *Lambda function ARN* in the context of intent *intent ARN*. Please check the resource-based policy on the function.

The response from the GetIntent operation contains a field called checksum that identifies a specific revision of the intent. When you use the [PutIntent](#) operation to update an intent, you must provide the checksum value. If you don't, you get the following error message:

An error occurred (PreconditionFailedException) when calling the PutIntent operation: Intent *intent name* already exists. If you are trying to update *intent name* you must specify the checksum.

This exercise uses the Lambda function from [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#). For instructions to create the Lambda function, see [Step 3: Create a Lambda Function \(Console\)](#).

Note

The following AWS CLI example is formatted for Unix, Linux, and macOS. For Windows, change "\$LATEST" to \$LATEST.

To add a Lambda function to an intent

1. In the AWS CLI, add the InvokeFunction permission for the OrderFlowers intent:

```
aws lambda add-permission \  
  --region region \  
  --function-name OrderFlowersCodeHook \  
  --statement-id LexGettingStarted-OrderFlowersBot \  
  --action lambda:InvokeFunction \  
  --principal lex.amazonaws.com \  
  --source-arn "arn:aws:lex:region:account ID:intent:OrderFlowers:*" \  
  --source-account account ID
```

Lambda sends the following response:

```
{
  "Statement": "{\\"Sid\\":\\"LexGettingStarted-OrderFlowersBot\\",
    \\"Resource\\":\\"arn:aws:lambda:region:account ID:function:OrderFlowersCodeHook
\\",
    \\"Effect\\":\\"Allow\\",
    \\"Principal\\":{\\"Service\\":\\"lex.amazonaws.com\\"},
    \\"Action\\":[\\"lambda:InvokeFunction\\"],
    \\"Condition\\":{\\"StringEquals\\":
      {\\"AWS:SourceAccount\\": \\"account ID\\",
      {\\"AWS:SourceArn\\":
        \\"arn:aws:lex:region:account ID:intent:OrderFlowers:*\\"}}}}"
}
```

2. Get the intent from Amazon Lex. Amazon Lex sends the output to a file called **OrderFlowers-V3.json**.

```
aws lex-models get-intent \
  --region region \
  --name OrderFlowers \
  --intent-version "\$LATEST" > OrderFlowers-V3.json
```

3. In a text editor, open the **OrderFlowers-V3.json**.

1. Find and delete the `createdDate`, `lastUpdatedDate`, and `version` fields.
2. Update the `fulfillmentActivity` field:

```
"fulfillmentActivity": {
  "type": "CodeHook",
  "codeHook": {
    "uri": "arn:aws:lambda:region:account ID:function:OrderFlowersCodeHook",
    "messageVersion": "1.0"
  }
}
```

3. Save the file.
4. In the AWS CLI, send the updated intent to Amazon Lex:

```
aws lex-models put-intent \
  --region region \
  --name OrderFlowers \
```

```
--cli-input-json file://OrderFlowers-V3.json
```

Now that you have updated the intent, rebuild the bot.

To rebuild the OrderFlowersBot bot

1. In the AWS CLI, get the definition of the OrderFlowersBot bot and save it to a file:

```
aws lex-models get-bot \  
  --region region \  
  --name OrderFlowersBot \  
  --version-or-alias "$LATEST" > OrderFlowersBot-V3.json
```

2. In a text editor, open **OrderFlowersBot-V3.json**. Remove the `createdDate`, `lastUpdatedDate`, `status`, and `version` fields.
3. In the text editor, add the following line to the definition of the bot:

```
"processBehavior": "BUILD",
```

4. In the AWS CLI, build a new revision of the bot:

```
aws lex-models put-bot \  
  --region region \  
  --name OrderFlowersBot \  
  --cli-input-json file://OrderFlowersBot-V3.json
```

The response from the server is:

```
{  
  "status": "READY",  
  "intents": [  
    {  
      "intentVersion": "$LATEST",  
      "intentName": "OrderFlowers"  
    }  
  ],  
  "name": "OrderFlowersBot",  
  "locale": "en-US",  
  "checksum": "checksum",  
  "abortStatement": {
```



```
    "messages": [
      {
        "content": "Sorry, I'm not able to assist at this time",
        "contentType": "PlainText"
      }
    ],
    "version": "$LATEST",
    "lastUpdatedDate": timestamp,
    "createdDate": timestamp,
    "clarificationPrompt": {
      "maxAttempts": 2,
      "messages": [
        {
          "content": "I didn't understand you, what would you like to do?",
          "contentType": "PlainText"
        }
      ]
    },
    "voiceId": "Salli",
    "childDirected": false,
    "idleSessionTTLInSeconds": 600,
    "description": "Bot to order flowers on the behalf of a user"
  }
}
```

Next Step

[Exercise 4: Publish a Version \(AWS CLI\)](#)

Exercise 4: Publish a Version (AWS CLI)

Now, create a version of the bot that you created in Exercise 1. A *version* is a snapshot of the bot. After you create a version, you can't change it. The only version of a bot that you can update is the \$LATEST version. For more information about versions, see [Versioning and Aliases](#).

Before you can publish a version of a bot, you must publish the intents that it uses. Likewise, you must publish the slot types that those intents refer to. In general, to publish a version of a bot, you do the following:

1. Publish a version of a slot type with the [CreateSlotTypeVersion](#) operation.
2. Publish a version of an intent with the [CreateIntentVersion](#) operation.

3. Publish a version of a bot with the [CreateBotVersion](#) operation .

To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#) .

Topics

- [Step 1: Publish the Slot Type \(AWS CLI\)](#)
- [Step 2: Publish the Intent \(AWS CLI\)](#)
- [Step 3: Publish the Bot \(AWS CLI\)](#)

Step 1: Publish the Slot Type (AWS CLI)

Before you can publish a version of any intents that use a slot type, you must publish a version of that slot type. In this case, you publish the `FlowerTypes` slot type.

Note

The following AWS CLI example is formatted for Unix, Linux, and macOS. For Windows, change `"\ $LATEST"` to `$LATEST` and replace the backslash (`\`) continuation character at the end of each line with a caret (`^`).

To publish a slot type (AWS CLI)

1. In the AWS CLI, get the latest version of the slot type:

```
aws lex-models get-slot-type \  
  --region region \  
  --name FlowerTypes \  
  --slot-type-version "\ $LATEST"
```

The response from Amazon Lex follows. Record the checksum for the current revision of the `$LATEST` version.

```
{  
  "enumerationValues": [  
    {  
      "value": "tulips"    }  
  ]  
}
```

```
    },
    {
      "value": "lilies"
    },
    {
      "value": "roses"
    }
  ],
  "name": "FlowerTypes",
  "checksum": "checksum",
  "version": "$LATEST",
  "lastUpdatedDate": timestamp,
  "createdDate": timestamp,
  "description": "Types of flowers to pick up"
}
```

2. Publish a version of the slot type. Use the checksum that you recorded in the previous step.

```
aws lex-models create-slot-type-version \
  --region region \
  --name FlowerTypes \
  --checksum "checksum"
```

The response from Amazon Lex follows. Record the version number for the next step.

```
{
  "version": "1",
  "enumerationValues": [
    {
      "value": "tulips"
    },
    {
      "value": "lilies"
    },
    {
      "value": "roses"
    }
  ],
  "name": "FlowerTypes",
  "createdDate": timestamp,
  "lastUpdatedDate": timestamp,
  "description": "Types of flowers to pick up"
}
```

```
}
```

Next Step

[Step 2: Publish the Intent \(AWS CLI\)](#)

Step 2: Publish the Intent (AWS CLI)

Before you can publish an intent, you have to publish all of the slot types referred to by the intent. The slot types must be numbered versions, not the \$LATEST version.

First, update the `OrderFlowers` intent to use the version of the `FlowerTypes` slot type that you published in the previous step. Then publish a new version of the `OrderFlowers` intent.

Note

The following AWS CLI example is formatted for Unix, Linux, and macOS. For Windows, change `"\ $LATEST"` to `$LATEST` and replace the backslash (`\`) continuation character at the end of each line with a caret (`^`).

To publish a version of an intent (AWS CLI)

1. In the AWS CLI, get the \$LATEST version of the `OrderFlowers` intent and save it to a file:

```
aws lex-models get-intent \  
  --region region \  
  --name OrderFlowers \  
  --intent-version "\ $LATEST" > OrderFlowers_V4.json
```

2. In a text editor, open the **OrderFlowers_V4.json** file. Delete the `createdDate`, `lastUpdatedDate`, and `version` fields. Find the `FlowerTypes` slot type and change the version to the version number that you recorded in the previous step. The following fragment of the **OrderFlowers_V4.json** file shows the location of the change:

```
{  
  "slotType": "FlowerTypes",  
  "name": "FlowerType",  
  "slotConstraint": "Required",  
  "valueElicitationPrompt": {
```

```
        "maxAttempts": 2,
        "messages": [
            {
                "content": "What type of flowers?",
                "contentType": "PlainText"
            }
        ]
    },
    "priority": 1,
    "slotTypeVersion": "version",
    "sampleUtterances": []
},
```

3. In the AWS CLI, save the revision of the intent:

```
aws lex-models put-intent \
    --name OrderFlowers \
    --cli-input-json file://OrderFlowers_V4.json
```

4. Get the checksum of the latest revision of the intent:

```
aws lex-models get-intent \
    --region region \
    --name OrderFlowers \
    --intent-version "$LATEST" > OrderFlowers_V4a.json
```

The following fragment of the response shows the checksum of the intent. Record this for the next step.

```
"name": "OrderFlowers",
"checksum": "checksum",
"version": "$LATEST",
```

5. Publish a new version of the intent:

```
aws lex-models create-intent-version \
    --region region \
    --name OrderFlowers \
    --checksum "checksum"
```

The following fragment of the response shows the new version of the intent. Record the version number for the next step.

```
"name": "OrderFlowers",  
"checksum": "checksum",  
"version": "version",
```

Next Step

[Step 3: Publish the Bot \(AWS CLI\)](#)

Step 3: Publish the Bot (AWS CLI)

After you have published all of the slot types and intents that are used by your bot, you can publish the bot.

Update the OrderFlowersBot bot to use the OrderFlowers intent that you updated in the previous step. Then, publish a new version of the OrderFlowersBot bot.

Note

The following AWS CLI example is formatted for Unix, Linux, and macOS. For Windows, change "\$LATEST" to \$LATEST and replace the backslash (\) continuation character at the end of each line with a caret (^).

To publish a version of a bot (AWS CLI)

1. In the AWS CLI, get the \$LATEST version of the OrderFlowersBot bot and save it to a file:

```
aws lex-models get-bot \  
  --region region \  
  --name OrderFlowersBot \  
  --version-or-alias "$LATEST" > OrderFlowersBot_V4.json
```

2. In a text editor, open the **OrderFlowersBot_V4.json** file. Delete the `createdDate`, `lastUpdatedDate`, `status` and `version` fields. Find the OrderFlowers intent and change the version to the version number that you recorded in the previous step. The following fragment of **OrderFlowersBot_V4.json** shows the location of the change.

```
"intents": [  
  {  
    "name": "OrderFlowers",  
    "checksum": "checksum",  
    "version": "version",
```

```
{
  "intentVersion": "version",
  "intentName": "OrderFlowers"
}
```

3. In the AWS CLI, save the new revision of the bot. Make note of the version number returned by the call to `put-bot`.

```
aws lex-models put-bot \
  --name OrderFlowersBot \
  --cli-input-json file://OrderFlowersBot_V4.json
```

4. Get the checksum of the latest revision of the bot. Use the version number returned in step 3.

```
aws lex-models get-bot \
  --region region \
  --version-or-alias version \
  --name OrderFlowersBot > OrderFlowersBot_V4a.json
```

The following fragment of the response shows the checksum of the bot. Record this for the next step.

```
"name": "OrderFlowersBot",
"locale": "en-US",
"checksum": "checksum",
```

5. Publish a new version of the bot:

```
aws lex-models create-bot-version \
  --region region \
  --name OrderFlowersBot \
  --checksum "checksum"
```

The following fragment of the response shows the new version of the bot.

```
"checksum": "checksum",
"abortStatement": {
  ...
},
"version": "1",
"lastUpdatedDate": timestamp,
```

Next Step

[Exercise 5: Create an Alias \(AWS CLI\)](#)

Exercise 5: Create an Alias (AWS CLI)

An alias is a pointer to a specific version of a bot. With an alias you can easily update the version that your client applications are using. For more information, see [Versioning and Aliases](#). To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#).

To create an alias (AWS CLI)

1. In the AWS CLI, get the version of the OrderFlowersBot bot that you created in [Exercise 4: Publish a Version \(AWS CLI\)](#).

```
aws lex-models get-bot \
  --region region \
  --name OrderFlowersBot \
  --version-or-alias version > OrderFlowersBot_V5.json
```

2. In a text editor, open **OrderFlowersBot_v5.json**. Find and record the version number.
3. In the AWS CLI, create the bot alias:

```
aws lex-models put-bot-alias \
  --region region \
  --name PROD \
  --bot-name OrderFlowersBot \
  --bot-version version
```

The following is the response from the server:

```
{
  "name": "PROD",
  "createdDate": timestamp,
  "checksum": "checksum",
  "lastUpdatedDate": timestamp,
  "botName": "OrderFlowersBot",
  "botVersion": "1"
}}
```


Next Step

[Exercise 6: Clean Up \(AWS CLI\)](#)

Exercise 6: Clean Up (AWS CLI)

Delete the resources that you created and clean up your account.

You can delete only resources that are not in use. In general, you should delete resources in the following order.

1. Delete aliases to free up bot resources.
2. Delete bots to free up intent resources.
3. Delete intents to free up slot type resources.
4. Delete slot types.

To run the commands in this exercise, you need to know the region where the commands will be run. For a list of regions, see [Model Building Quotas](#).

To clean up your account (AWS CLI)

1. In the AWS CLI command line, delete the alias:

```
aws lex-models delete-bot-alias \  
  --region region \  
  --name PROD \  
  --bot-name OrderFlowersBot
```

2. In the AWS CLI command line, delete the bot:

```
aws lex-models delete-bot \  
  --region region \  
  --name OrderFlowersBot
```

3. In the AWS CLI command line, delete the intent:

```
aws lex-models delete-intent \  
  --region region \  
  --name OrderFlowers
```

4. From the AWS CLI command line, delete the slot type:

```
aws lex-models delete-slot-type \  
  --region region \  
  --name FlowerTypes
```

You have removed all of the resources that you created and cleaned up your account.

Versioning and Aliases

Amazon Lex supports publishing versions of bots, intents, and slot types so that you can control the implementation that your client applications use. A *version* is a numbered snapshot of your work that you can publish for use in different parts of your workflow, such as development, beta deployment, and production.

Amazon Lex bots also support aliases. An *alias* is a pointer to a specific version of a bot. With an alias, you can easily update the version that your client applications are using. For example, you can point an alias to version 1 of your bot. When you are ready to update the bot, you publish version 2 and change the alias to point to the new version. Because your applications use the alias instead of a specific version, all of your clients get the new functionality without needing to be updated.

Topics

- [Versioning](#)
- [Aliases](#)

Versioning

When you version an Amazon Lex resource you create a snapshot of the resource so that you can use the resource as it existed when the version was made. Once you've created a version it will stay the same while you continue to work on your application.

The \$LATEST Version

When you create an Amazon Lex bot, intent, or slot type there is only one version, the \$LATEST version.




\$LATEST is the working copy of your resource. You can update only the \$LATEST version and until you publish your first version, \$LATEST is the only version of the resource that you have.

Only the \$LATEST version of a resource can use the \$LATEST version of another resource. For example, the \$LATEST version of a bot can use the \$LATEST version of an intent, and the \$LATEST version of an intent can use the \$LATEST version of a slot type.

The \$LATEST version of your bot should only be used for manual testing. Amazon Lex limits the number of runtime requests that you can make to the \$LATEST version of the bot.

Publishing an Amazon Lex Resource Version

When you publish a resource, Amazon Lex makes a copy of the \$LATEST version and saves it as a numbered version. The published version can't be changed.



You create and publish versions using the Amazon Lex console or the [CreateBotVersion](#) operation. For an example, see [Exercise 3: Publish a Version and Create an Alias](#).

When you modify the \$LATEST version of a resource, you can publish the new version to make the changes available to your client applications. Every time you publish a version, Amazon Lex copies the \$LATEST version to create the new version and increments the version number by 1. Version numbers are never reused. For example, if you remove a resource numbered version 10 and then recreate it, the next version number Amazon Lex assigns is version 11.

Before you can publish a bot, you must point it to a numbered version of any intent that it uses. If you try to publish a new version of a bot that uses the \$LATEST version of an intent, Amazon Lex returns an HTTP 400 Bad Request exception. Before you can publish a numbered version of the intent, you must point the intent to a numbered version of any slot type that it uses. Otherwise you will get an HTTP 400 Bad Request exception.



Note

Amazon Lex publishes a new version only if the last published version is different from the \$LATEST version. If you try to publish the \$LATEST version without modifying it, Amazon Lex doesn't create or publish a new version.

Updating an Amazon Lex Resource

You can update only the \$LATEST version of an Amazon Lex bot, intent, or slot type. Published versions can't be changed. You can publish a new version any time after you update a resource in the console or with the [CreateBotVersion](#), the [CreateIntentVersion](#) or the [CreateSlotTypeVersion](#) operations.

Deleting an Amazon Lex Resource or Version

Amazon Lex supports deleting a resource or version using the console or one of the API operations:

- [DeleteBot](#)
- [DeleteBotVersion](#)
- [DeleteBotAlias](#)
- [DeleteBotChannelAssociation](#)
- [DeleteIntent](#)
- [DeleteIntentVersion](#)
- [DeleteSlotType](#)
- [DeleteSlotTypeVersion](#)


Aliases

An alias is a pointer to a specific version of an Amazon Lex bot. Use an alias to allow client applications to use a specific version of the bot without requiring the application to track which version that is.

The following example shows two versions of an Amazon Lex bot, version version 1 and version 2. Each of these bot versions has an associated alias, BETA and PROD, respectively. Client applications use the PROD alias to access the bot.



When you create a second version of the bot, you can update the alias to point to the new version of the bot using the console or the [PutBot](#) operation. When you change the alias, all of your client applications use the new version. If there is a problem with the new version, you can roll back to the previous version by simply changing the alias to point to that version.



Note

Although you can test the `$LATEST` version of a bot in the console, we recommend that when you integrate a bot with your client application, you first publish a version and create an alias that points to that version. Use the alias in your client application for the reasons explained in this section. When you update an alias, Amazon Lex will wait until the session timeout of all current sessions expires before it starts using the new version. For more information about the session timeout, see [the section called “Setting the Session Timeout”](#)

Using Lambda Functions

You can create AWS Lambda functions to use as code hooks for your Amazon Lex bot. You can identify Lambda functions to perform initialization and validation, fulfillment, or both in your intent configuration.

We recommend that you use a Lambda function as a code hook for your bot. Without a Lambda function, your bot returns the intent information to the client application for fulfillment.

Topics

- [Lambda Function Input Event and Response Format](#)
- [Amazon Lex and AWS Lambda Blueprints](#)

Lambda Function Input Event and Response Format

This section describes the structure of the event data that Amazon Lex provides to a Lambda function. Use this information to parse the input in your Lambda code. It also explains the format of the response that Amazon Lex expects your Lambda function to return.

Topics

- [Input Event Format](#)
- [Response Format](#)

Input Event Format

The following shows the general format of an Amazon Lex event that is passed to a Lambda function. Use this information when you are writing your Lambda function.

Note

The input format may change without a corresponding change in the `messageVersion`. Your code should not throw an error if new fields are present.

```
{  
  "currentIntent": {
```



```
"name": "intent-name",
"nluIntentConfidenceScore": score,
"slots": {
  "slot name": "value",
  "slot name": "value"
},
"slotDetails": {
  "slot name": {
    "resolutions" : [
      { "value": "resolved value" },
      { "value": "resolved value" }
    ],
    "originalValue": "original text"
  },
  "slot name": {
    "resolutions" : [
      { "value": "resolved value" },
      { "value": "resolved value" }
    ],
    "originalValue": "original text"
  }
},
"confirmationStatus": "None, Confirmed, or Denied (intent confirmation, if configured)"
},
"alternativeIntents": [
  {
    "name": "intent-name",
    "nluIntentConfidenceScore": score,
    "slots": {
      "slot name": "value",
      "slot name": "value"
    },
    "slotDetails": {
      "slot name": {
        "resolutions" : [
          { "value": "resolved value" },
          { "value": "resolved value" }
        ],
        "originalValue": "original text"
      },
      "slot name": {
        "resolutions" : [
          { "value": "resolved value" },
```

```

        { "value": "resolved value" }
      ],
      "originalValue": "original text"
    }
  },
  "confirmationStatus": "None, Confirmed, or Denied (intent confirmation, if configured)"
},
"bot": {
  "name": "bot name",
  "alias": "bot alias",
  "version": "bot version"
},
"userId": "User ID specified in the POST request to Amazon Lex.",
"inputTranscript": "Text used to process the request",
"invocationSource": "FulfillmentCodeHook or DialogCodeHook",
"outputDialogMode": "Text or Voice, based on ContentType request header in runtime API request",
"messageVersion": "1.0",
"sessionAttributes": {
  "key": "value",
  "key": "value"
},
"requestAttributes": {
  "key": "value",
  "key": "value"
},
"recentIntentSummaryView": [
  {
    "intentName": "Name",
    "checkpointLabel": Label,
    "slots": {
      "slot name": "value",
      "slot name": "value"
    },
    "confirmationStatus": "None, Confirmed, or Denied (intent confirmation, if configured)",
    "dialogActionType": "ElicitIntent, ElicitSlot, ConfirmIntent, Delegate, or Close",
    "fulfillmentState": "Fulfilled or Failed",
    "slotToElicit": "Next slot to elicit"
  }
],

```

```
"sentimentResponse": {
  "sentimentLabel": "sentiment",
  "sentimentScore": "score"
},
"kendraResponse": {
  Complete query response from Amazon Kendra
},
"activeContexts": [
  {
    "timeToLive": {
      "timeToLiveInSeconds": seconds,
      "turnsToLive": turns
    },
    "name": "name",
    "parameters": {
      "key name": "value"
    }
  }
]
```

Note the following additional information about the event fields:

- **currentIntent** – Provides the intent name, slots, slotDetails and confirmationStatus fields.

nluIntentConfidenceScore is the confidence that Amazon Lex has that the current intent is the one that best matches the user's current intent.

slots is a map of slot names, configured for the intent, to slot values that Amazon Lex has recognized in the user conversation. A slot value remains null until the user provides a value.

The slot value in the input event may not match one of the values configured for the slot. For example, if the user responds to the prompt "What color car would you like?" with "pizza," Amazon Lex will return "pizza" as the slot value. Your function should validate the values to make sure that they make sense in context.

`slotDetails` provides additional information about a slot value. The `resolutions` array contains a list of additional values recognized for the slot. Each slot can have a maximum of five values.

The `originalValue` field contains the value that was entered by the user for the slot. When the slot type is configured to return the top resolution value as the slot value, the `originalValue` may be different from the value in the `slots` field.

`confirmationStatus` provides the user response to a confirmation prompt, if there is one. For example, if Amazon Lex asks "Do you want to order a large cheese pizza?," depending on the user response, the value of this field can be `Confirmed` or `Denied`. Otherwise, this value of this field is `None`.

If the user confirms the intent, Amazon Lex sets this field to `Confirmed`. If the user denies the intent, Amazon Lex sets this value to `Denied`.

In the confirmation response, a user utterance might provide slot updates. For example, the user might say "yes, change size to medium." In this case, the subsequent Lambda event has the updated slot value, `PizzaSize` set to `medium`. Amazon Lex sets the `confirmationStatus` to `None`, because the user modified some slot data, requiring the Lambda function to perform user data validation.

- **alternativeIntents** – If you enable confidence scores, Amazon Lex returns up to four alternative intents. Each intent includes a score that indicates the level of confidence that Amazon Lex has that the intent is the correct intent based on the user's utterance.

The contents of the alternative intents is the same as the contents of the `currentIntent` field. For more information, see [Using Confidence Scores](#).

- **bot** – Information about the bot that processed the request.
 - `name` – The name of the bot that processed the request.
 - `alias` – The alias of the bot version that processed the request.
 - `version` – The version of the bot that processed the request.
- **userId** – This value is provided by the client application. Amazon Lex passes it to the Lambda function.
- **inputTranscript** – The text used to process the request.

If the input was text, the `inputTranscript` field contains the text that was input by the user.

If the input was an audio stream, the `inputTranscript` field contains the text extracted from the audio stream. This is the text that is actually processed to recognize intents and slot values.

- **invocationSource** – To indicate why Amazon Lex is invoking the Lambda function, it sets this to one of the following values:
 - `DialogCodeHook` – Amazon Lex sets this value to direct the Lambda function to initialize the function and to validate the user's data input.

When the intent is configured to invoke a Lambda function as an initialization and validation code hook, Amazon Lex invokes the specified Lambda function on each user input (utterance) after Amazon Lex understands the intent.

Note

If the intent is not clear, Amazon Lex can't invoke the Lambda function.

- **FulfillmentCodeHook** – Amazon Lex sets this value to direct the Lambda function to fulfill an intent.

If the intent is configured to invoke a Lambda function as a fulfillment code hook, Amazon Lex sets the `invocationSource` to this value only after it has all the slot data to fulfill the intent.

In your intent configuration, you can have two separate Lambda functions to initialize and validate user data and to fulfill the intent. You can also use one Lambda function to do both. In that case, your Lambda function can use the `invocationSource` value to follow the correct code path.

- **outputDialogMode** – For each user input, the client sends the request to Amazon Lex using one of the runtime API operations, [PostContent](#) or [PostText](#). Amazon Lex uses the request parameters to determine whether the response to the client is text or voice, and sets this field accordingly.

The Lambda function can use this information to generate an appropriate message. For example, if the client expects a voice response, your Lambda function could return Speech Synthesis Markup Language (SSML) instead of text.

- **messageVersion** – The version of the message that identifies the format of the event data going into the Lambda function and the expected format of the response from a Lambda function.

Note

You configure this value when you define an intent. In the current implementation, only message version 1.0 is supported. Therefore, the console assumes the default value of 1.0 and doesn't show the message version.

- **sessionAttributes** – Application-specific session attributes that the client sends in the request. If you want Amazon Lex to include them in the response to the client, your Lambda function should send these back to Amazon Lex in the response. For more information, see [Setting Session Attributes](#)
- **requestAttributes** – Request-specific attributes that the client sends in the request. Use request attributes to pass information that doesn't need to persist for the entire session. If there are no request attributes, the value will be null. For more information, see [Setting Request Attributes](#)
- **recentIntentSummaryView** – Information about the state of an intent. You can see information about the last three intents used. You can use this information to set values in the intent or to return to a previous intent. For more information, see [Managing Sessions With the Amazon Lex API](#).
- **sentimentResponse** – The result of an Amazon Comprehend sentiment analysis of the last utterance. You can use this information to manage the conversation flow of your bot depending on the sentiment expressed by the user. For more information, see [Sentiment Analysis](#).
- **kendraResponse** – The result of a query to an Amazon Kendra index. Only present in the input to a fulfillment code hook and only when the intent extends the `AMAZON.KendraSearchIntent` built-in intent. The field contains the entire response from the Amazon Kendra search. For more information, see [AMAZON.KendraSearchIntent](#).
- **activeContexts** – One or more contexts that are active during this turn of a conversation with the user.

- **timeToLive** – The length of time or number of turns in the conversation with the user that the context remains active.
- **name** – the name of the context.
- **parameters** a list of key/value pairs that contains the name and value of the slots from the intent that activated the context.

For more information, see [Setting Intent Context](#).

Response Format

Amazon Lex expects a response from a Lambda function in the following format:

```
{
  "sessionAttributes": {
    "key1": "value1",
    "key2": "value2"
    ...
  },
  "recentIntentSummaryView": [
    {
      "intentName": "Name",
      "checkpointLabel": "Label",
      "slots": {
        "slot name": "value",
        "slot name": "value"
      },
      "confirmationStatus": "None, Confirmed, or Denied (intent confirmation, if configured)",
      "dialogActionType": "ElicitIntent, ElicitSlot, ConfirmIntent, Delegate, or Close",
      "fulfillmentState": "Fulfilled or Failed",
      "slotToElicit": "Next slot to elicit"
    }
  ],
  "activeContexts": [
    {
      "timeToLive": {
        "timeToLiveInSeconds": seconds,
        "turnsToLive": turns
      },
      "name": "name",
```



```
    "parameters": {
      "key name": "value"
    }
  ],
  "dialogAction": {
    "type": "ElicitIntent, ElicitSlot, ConfirmIntent, Delegate, or Close",
    Full structure based on the type field. See below for details.
  }
}
```

The response consists of four fields. The `sessionAttributes`, `recentIntentSummaryView`, and `activeContexts` fields are optional, the `dialogAction` field is required. The contents of the `dialogAction` field depends on the value of the `type` field. For details, see [dialogAction](#).

sessionAttributes

Optional. If you include the `sessionAttributes` field it can be empty. If your Lambda function doesn't return session attributes, the last known `sessionAttributes` passed via the API or Lambda function remain. For more information, see the [PostContent](#) and [PostText](#) operations.

```
"sessionAttributes": {
  "key1": "value1",
  "key2": "value2"
}
```

recentIntentSummaryView

Optional. If included, sets values for one or more recent intents. You can include information for up to three intents. For example, you can set values for previous intents based on information gathered by the current intent. The information in the summary must be valid for the intent. For example, the intent name must be an intent in the bot. If you include a slot value in the summary view, the slot must exist in the intent. If you don't include the `recentIntentSummaryView` in your response, all of the values for the recent intents remain unchanged. For more information, see the [PutSession](#) operation or the [IntentSummary](#) data type.

```
"recentIntentSummaryView": [
  {
    "intentName": "Name",
    "checkpointLabel": "Label",
    "slots": {
```

```
        "slot name": "value",
        "slot name": "value"
    },
    "confirmationStatus": "None, Confirmed, or Denied (intent confirmation, if configured)",
    "dialogActionType": "ElicitIntent, ElicitSlot, ConfirmIntent, Delegate, or Close",
    "fulfillmentState": "Fulfilled or Failed",
    "slotToElicit": "Next slot to elicit"
}
]
```

activeContexts

Optional. If included, sets the value for one or more contexts. For example, you can include a context to make one or more intents that have that context as an input eligible for recognition in the next turn of the conversation.

Any active contexts that are not included in the response have their time-to-live values decremented and may still be active on the next request.

If you specify a time-to-live of 0 for a context that was included in the input event, it will be inactive on the next request.

For more information, see [Setting Intent Context](#).

dialogAction

Required. The dialogAction field directs Amazon Lex to the next course of action, and describes what to expect from the user after Amazon Lex returns a response to the client.

The type field indicates the next course of action. It also determines the other fields that the Lambda function needs to provide as part of the dialogAction value.

- Close — Informs Amazon Lex not to expect a response from the user. For example, "Your pizza order has been placed" does not require a response.

The fulfillmentState field is required. Amazon Lex uses this value to set the dialogState field in the [PostContent](#) or [PostText](#) response to the client application. The message and

responseCard fields are optional. If you don't specify a message, Amazon Lex uses the goodbye message or the follow-up message configured for the intent.

```
"dialogAction": {
  "type": "Close",
  "fulfillmentState": "Fulfilled or Failed",
  "message": {
    "contentType": "PlainText or SSML or CustomPayload",
    "content": "Message to convey to the user. For example, Thanks, your pizza has been ordered."
  },
  "responseCard": {
    "version": integer-value,
    "contentType": "application/vnd.amazonaws.card.generic",
    "genericAttachments": [
      {
        "title": "card-title",
        "subTitle": "card-sub-title",
        "imageUrl": "URL of the image to be shown",
        "attachmentLinkUrl": "URL of the attachment to be associated with the card",
        "buttons": [
          {
            "text": "button-text",
            "value": "Value sent to server on button click"
          }
        ]
      }
    ]
  }
}
```

- **ConfirmIntent** — Informs Amazon Lex that the user is expected to give a yes or no answer to confirm or deny the current intent.

You must include the `intentName` and `slots` fields. The `slots` field must contain an entry for each of the filled slots for the specified intent. You don't need to include a entry in the `slots` field for slots that aren't filled. You must include the `message` field if the intent's `confirmationPrompt` field is null. The contents of the `message` field returned by the

Lambda function take precedence over the confirmationPrompt specified in the intent. The responseCard field is optional.

```
"dialogAction": {
  "type": "ConfirmIntent",
  "message": {
    "contentType": "PlainText or SSML or CustomPayload",
    "content": "Message to convey to the user. For example, Are you sure you want a
large pizza?"
  },
  "intentName": "intent-name",
  "slots": {
    "slot-name": "value",
    "slot-name": "value",
    "slot-name": "value"
  },
  "responseCard": {
    "version": integer-value,
    "contentType": "application/vnd.amazonaws.card.generic",
    "genericAttachments": [
      {
        "title": "card-title",
        "subTitle": "card-sub-title",
        "imageUrl": "URL of the image to be shown",
        "attachmentLinkUrl": "URL of the attachment to be associated with the
card",
        "buttons": [
          {
            "text": "button-text",
            "value": "Value sent to server on button click"
          }
        ]
      }
    ]
  }
}
```

- **Delegate** — Directs Amazon Lex to choose the next course of action based on the bot configuration. If the response does not include any session attributes Amazon Lex retains the existing attributes. If you want a slot value to be null, you don't need to include the slot field

in the request. You will get a `DependencyFailedException` exception if your fulfillment function returns the `Delegate` dialog action without removing any slots.

The `kendraQueryRequestPayload` and `kendraQueryFilterString` fields are optional and only used when the intent is derived from the `AMAZON.KendraSearchIntent` built-in intent. for more information, see [AMAZON.KendraSearchIntent](#).

```
"dialogAction": {
  "type": "Delegate",
  "slots": {
    "slot-name": "value",
    "slot-name": "value",
    "slot-name": "value"
  },
  "kendraQueryRequestPayload": "Amazon Kendra query",
  "kendraQueryFilterString": "Amazon Kendra attribute filters"
}
```

- **ElicitIntent** — Informs Amazon Lex that the user is expected to respond with an utterance that includes an intent. For example, "I want a large pizza," which indicates the `OrderPizzaIntent`. The utterance "large," on the other hand, is not sufficient for Amazon Lex to infer the user's intent.

The `message` and `responseCard` fields are optional. If you don't provide a message, Amazon Lex uses one of the bot's clarification prompts. If there is no clarification prompt defined, Amazon Lex returns a 400 Bad Request exception.

```
{
  "dialogAction": {
    "type": "ElicitIntent",
    "message": {
      "contentType": "PlainText or SSML or CustomPayload",
      "content": "Message to convey to the user. For example, What can I help you with?"
    },
    "responseCard": {
      "version": integer-value,
      "contentType": "application/vnd.amazonaws.card.generic",
      "genericAttachments": [
        {
```

```

        "title": "card-title",
        "subTitle": "card-sub-title",
        "imageUrl": "URL of the image to be shown",
        "attachmentLinkUrl": "URL of the attachment to be associated with the
card",
        "buttons": [
            {
                "text": "button-text",
                "value": "Value sent to server on button click"
            }
        ]
    }
}

```

- **ElicitSlot** — Informs Amazon Lex that the user is expected to provide a slot value in the response.

The `intentName`, `slotToElicit`, and `slots` fields are required. The `message` and `responseCard` fields are optional. If you don't specify a message, Amazon Lex uses one of the slot elicitation prompts configured for the slot.

```

"dialogAction": {
    "type": "ElicitSlot",
    "message": {
        "contentType": "PlainText or SSML or CustomPayload",
        "content": "Message to convey to the user. For example, What size pizza would
you like?"
    },
    "intentName": "intent-name",
    "slots": {
        "slot-name": "value",
        "slot-name": "value",
        "slot-name": "value"
    },
    "slotToElicit": "slot-name",
    "responseCard": {
        "version": integer-value,
        "contentType": "application/vnd.amazonaws.card.generic",

```

```
    "genericAttachments": [
      {
        "title": "card-title",
        "subTitle": "card-sub-title",
        "imageUrl": "URL of the image to be shown",
        "attachmentLinkUrl": "URL of the attachment to be associated with the
card",
        "buttons": [
          {
            "text": "button-text",
            "value": "Value sent to server on button click"
          }
        ]
      }
    ]
  }
}
```

Amazon Lex and AWS Lambda Blueprints

The Amazon Lex console provides example bots (called bot blueprints) that are preconfigured so you can quickly create and test a bot in the console. For each of these bot blueprints, Lambda function blueprints are also provided. These blueprints provide sample code that works with their corresponding bots. You can use these blueprints to quickly create a bot that is configured with a Lambda function as a code hook, and test the end-to-end setup without having to write code.

You can use the following Amazon Lex bot blueprints and the corresponding AWS Lambda function blueprints as code hooks for bots:

- Amazon Lex blueprint — OrderFlowers
 - AWS Lambda blueprint — `lex-order-flowers-python`
- Amazon Lex blueprint — ScheduleAppointment
 - AWS Lambda blueprint — `lex-make-appointment-python`
- Amazon Lex blueprint — BookTrip
 - AWS Lambda blueprint — `lex-book-trip-python`

To create a bot using a blueprint and configure it to use a Lambda function as a code hook, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#). For an example of using other blueprints, see [Additional Examples: Creating Amazon Lex Bots](#).

Updating a Blueprint for a Specific Locale

If you are using a blueprint in a locale other than English (US) (en-US), you need to update the name of any intents to include the locale. For example, if you are using the `OrderFlowers` blueprint, you need to do the following.

- Find the dispatch function near the end of the Lambda function code.
- In the dispatch function, update the name of the intent to include the locale that you are using. For example, if you are using the English (Australian) (en-AU) locale, change the line:

```
if intent_name == 'OrderFlowers':  
  
    to  
  
    if intent_name == 'OrderFlowers_enAU':
```

Other blueprints use other intent names, they should be updated as above before you use them.

Deploying Amazon Lex Bots

This section provides examples of deploying Amazon Lex bots on various messaging platforms and in mobile applications.

Topics

- [Deploying an Amazon Lex Bot on a Messaging Platform](#)
- [Deploying an Amazon Lex Bot in Mobile Applications](#)

Deploying an Amazon Lex Bot on a Messaging Platform

This section explains how to deploy Amazon Lex bots on the Facebook, Slack, and Twilio messaging platforms.

Note

When storing your Facebook, Slack, or Twilio configurations, Amazon Lex uses AWS Key Management Service customer managed keys to encrypt the information. The first time that you create a channel to one of these messaging platforms, Amazon Lex creates a default customer managed key (`aws/lex`). Alternatively, you can create your own customer managed key with AWS KMS. This gives you more flexibility, including the ability to create, rotate, and disable keys. You can also define access controls and audit the encryption keys used to protect your data. For more information, see the [AWS Key Management Service Developer Guide](#).

When a messaging platform sends a request to Amazon Lex it includes platform-specific information as a request attribute to your Lambda function. Use these attributes to customize the way that your bot behaves. For more information, see [Setting Request Attributes](#).

All of the attributes take the namespace, `x-amz-lex:`, as the prefix. For example, the `user-id` attribute is called `x-amz-lex:user-id`. There are common attributes that are sent by all messaging platforms in addition to attributes that are specific to a particular platform. The following tables list the request attributes that messaging platforms send to your bot's Lambda function.

Common Request Attributes

Attribute	Description
channel-id	The channel endpoint identifier from Amazon Lex.
channel-name	The channel name from Amazon Lex.
channel-type	One of the following values: • Facebook • Kik • Slack • Twilio-SMS
webhook-endpoint-url	The Amazon Lex endpoint for the channel.

Facebook Request Attributes

Attribute	Description
user-id	The Facebook identifier of the sender. See https://developers.facebook.com/docs/messenger-platform/webhook-reference/message-received .
facebook-page-id	The Facebook page identifier of the recipient. See https://developers.facebook.com/docs/messenger-platform/webhook-reference/message-received .

Kik Request Attributes

Attribute	Description
kik-chat-id	The identifier for the conversation that your bot is involved in. For more information, see https://dev.kik.com/#/docs/messaging#message-formats .

Attribute	Description
kik-chat-type	The type of conversation that the message originated from. For more information, see https://dev.kik.com/#/docs/messaging#message-formats .
kik-message-id	A UUID that identifies the message. For more information, see https://dev.kik.com/#/docs/messaging#message-formats .
kik-message-type	The type of message. For more information, see https://dev.kik.com/#/docs/messaging#message-types .

Twilio Request Attributes

Attribute	Description
user-id	The sender's phone number ("From"). See https://www.twilio.com/docs/api/rest/message .
twilio-target-phone-number	The phone number of the recipient ("To"). See https://www.twilio.com/docs/api/rest/message .

Slack Request Attributes

Attribute	Description
user-id	The Slack user identifier. See https://api.slack.com/types/user .
slack-team-id	The identifier of the team that sent the message. See https://api.slack.com/methods/team.info .
slack-bot-token	The developer token that gives the bot access to the Slack APIs. See https://api.slack.com/docs/token-types .

Integrating an Amazon Lex Bot with Facebook Messenger

This exercise shows how to integrate Facebook Messenger with your Amazon Lex bot. You perform the following steps:

1. Create an Amazon Lex bot
2. Create a Facebook application
3. Integrate Facebook Messenger with your Amazon Lex bot
4. Validate the integration

Topics

- [Step 1: Create an Amazon Lex Bot](#)
- [Step 2: Create a Facebook Application](#)
- [Step 3: Integrate Facebook Messenger with the Amazon Lex Bot](#)
- [Step 4: Test the Integration](#)

Step 1: Create an Amazon Lex Bot

If you don't already have an Amazon Lex bot, create and deploy one. In this topic, we assume that you are using the bot that you created in Getting Started Exercise 1. However, you can use any of the example bots provided in this guide. For Getting Started Exercise 1, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#).

1. Create an Amazon Lex bot. For instructions, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#).
2. Deploy the bot and create an alias. For instructions, see [Exercise 3: Publish a Version and Create an Alias](#).

Step 2: Create a Facebook Application

On the Facebook developer portal, create a Facebook application and a Facebook page. For instructions, see [Quick Start](#) in the Facebook Messenger platform documentation. Write down the following:

- The **App Secret** for the Facebook App

- The **Page Access Token** for the Facebook page

Step 3: Integrate Facebook Messenger with the Amazon Lex Bot

In this section, you integrate Facebook Messenger with your Amazon Lex bot.

After you complete this step, the console provides a callback URL. Write down this URL.

To integrate Facebook Messenger with your bot

1. a. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
b. Choose your Amazon Lex bot.
c. Choose **Channels**.
d. Choose **Facebook** under **Chatbots**. The console displays the Facebook integration page.
e. On the Facebook integration page, do the following:
 - Type the following name: BotFacebookAssociation.
 - For **KMS key**, choose **aws/lex**.
 - For **Alias**, choose the bot alias.
 - For **Verify token**, type a token. This can be any string you choose (for example, ExampleToken). You use this token later in the Facebook developer portal when you set up the webhook.
 - For **Page access token**, type the token that you obtained from Facebook in Step 2.
 - For **App secret key**, type the key that you obtained from Facebook in Step 2.

The screenshot shows the Amazon Lex console interface for configuring a bot channel. The bot is named 'BookTrip' and is in the 'Latest' state. The 'Channels' tab is active, showing the configuration for the 'Facebook' channel. The form includes the following fields:

- Name:** BotFacebookAssociation
- Description:** Channel for associating Facebook
- IAM Role:** AWSServiceRoleForLexChannels (Automatically created on your behalf)
- KMS key:** aws/lex
- Alias:** Beta
- Verify token:** ExampleToken
- Page access token:** Page access token
- App secret key:** App secret key

An **Activate** button is located at the bottom of the form. A **Test Bot** button is in the bottom right corner.

f. Choose **Activate**.

The console creates the bot channel association and returns a callback URL. Write down this URL.

- On the Facebook developer portal, choose your app.
- Choose the **Messenger** product, and choose **Setup webhooks** in the **Webhooks** section of the page.

For instructions, see [Quick Start](#) in the Facebook Messenger platform documentation.

- On the **webhook** page of the subscription wizard, do the following:
 - For **Callback URL**, type the callback URL provided in the Amazon Lex console earlier in the procedure.
 - For **Verify Token**, type the same token that you used in Amazon Lex.
 - Choose **Subscription Fields** (**messages**, **messaging_postbacks**, and **messaging_optins**).
 - Choose **Verify and Save**. This initiates a handshake between Facebook and Amazon Lex.
- Enable Webhooks integration. Choose the page that you created, and then choose **subscribe**.

Note

If you update or recreate a webhook, unsubscribe and then resubscribe to the page.

Step 4: Test the Integration

You can now start a conversation from Facebook Messenger with your Amazon Lex bot.

1. Open your Facebook page, and choose **Message**.
2. In the Messenger window, use the same test utterances provided in [Step 1: Create an Amazon Lex Bot \(Console\)](#).

Integrating an Amazon Lex Bot with Kik

This exercise provides instructions for integrating an Amazon Lex bot with the Kik messaging application. You perform the following steps:

1. Create an Amazon Lex bot.
2. Create a Kik bot using the Kik app and website.
3. Integrate the your Amazon Lex bot with the Kik bot using the Amazon Lex console.
4. Engage in a conversation with your Amazon Lex bot using Kik to test the association between your Amazon Lex bot and Kik.

Topics

- [Step 1: Create an Amazon Lex Bot](#)
- [Step 2: Create a Kik Bot](#)
- [Step 3: Integrate the Kik Bot with the Amazon Lex Bot](#)
- [Step 4: Test the Integration](#)

Step 1: Create an Amazon Lex Bot

If you don't already have an Amazon Lex bot, create and deploy one. In this topic, we assume that you are using the bot that you created in Getting Started Exercise 1. However, you can use any of

the example bots provided in this guide. For Getting Started Exercise 1, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#)

1. Create an Amazon Lex bot. For instructions, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#).
2. Deploy the bot and create an alias. For instructions, see [Exercise 3: Publish a Version and Create an Alias](#).

Next Step

[Step 2: Create a Kik Bot](#)

Step 2: Create a Kik Bot

In this step you use the Kik user interface to create a Kik bot. You use information generated while creating the bot to connect it to your Amazon Lex bot.

1. If you haven't already, download and install the Kik app and sign up for a Kik account. If you have an account, log in.
2. Open the Kik website at <https://dev.kik.com/>. Leave the browser window open.
3. In the Kik app, choose the gear icon to open settings, and then choose **Your Kik Code**.
4. Scan the Kik code on the Kik website to open the Botsworth chatbot. Choose **Yes** to open the Bot Dashboard.
5. In the Kik app, choose **Create a Bot**. Follow the prompts to create your Kik bot.
6. Once the bot is created, choose **Configuration** in your browser. Make sure that your new bot is selected.
7. Note the bot name and the API key for the next section.

Next Step

[Step 3: Integrate the Kik Bot with the Amazon Lex Bot](#)

Step 3: Integrate the Kik Bot with the Amazon Lex Bot

Now that you have created an Amazon Lex bot and a Kik bot, you are ready to create an channel association between them in Amazon Lex. When the association is activated, Amazon Lex automatically sets up a callback URL with Kik.

1. Sign in to the AWS Management Console, and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose the Amazon Lex bot that you created in Step 1.
3. Choose the **Channels** tab.
4. In the **Channels** section, choose **Kik**.
5. On the Kik page, provide the following:
 - Type a name. For example, BotKikIntegration.
 - Type a description.
 - Choose "aws/lex" from the **KMS key** drop-down.
 - For **Alias**, choose an alias from the drop-down.
 - For **Kik bot user name**, type the name that you gave the bot on Kik.
 - For **Kik API key**, type the API key that was assigned to the bot on Kik.
 - For **User greeting**, type the greeting that you would like your bot to send the first time that a user chats with it.
 - For **Error message**, enter an error message that is shown to the user when part of the conversation is not understood.
 - For **Group chat behavior**, choose one of the options:
 - **Enable** – Enables the entire chat group to interact with your bot in a single conversation.
 - **Disable** – Restricts the conversation to one user in the chat group.
 - Choose **Activate** to create the association and link it to the Kik bot.

Kik

Fill in the form below and click activate to get a callback URL to use with Kik. You can generate multiple callback URLs. [Learn more](#) on steps to integrate with Kik.

Channel Name*	KikBotIntegration	
Channel Description	Integrate an Amazon Lex bot with Kik	
IAM Role	AWSServiceRoleForLexChannels Automatically created on your behalf	
KMS key	aws/lex	
Alias*	BETA	
Kik Bot User Name*	XXXXXXXX	
Kik API Key*	XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXX	
User Greeting*	Welcome to my first Amazon Lex bot on Kik	

Advanced configuration

Error Message*	There seems to be a problem.	
Group Chat Behavior	☐ Enable ☒ Disable	

* Required Field


Next Step

[Step 4: Test the Integration](#)

Step 4: Test the Integration

Now that you have created an association between your Amazon Lex bot and Kik, you can use the Kik app to test the association.

1. Start the Kik app and log in. Select the bot that you created.
2. You can test the bot with the following:



As you enter each phrase, your Amazon Lex bot will respond through Kik with the prompt that you created for each slot.

Integrating an Amazon Lex Bot with Slack

This exercise provides instructions for integrating an Amazon Lex bot with the Slack messaging application. You perform the following steps:

1. Create an Amazon Lex bot.
2. Create a Slack messaging application.
3. Integrate the Slack application with your bot Amazon Lex.

4. Test the integration by engaging in conversation with your Amazon Lex bot. You send messages with the Slack application and test in a browser window.

Topics

- [Step 1: Create an Amazon Lex Bot](#)
- [Step 2: Sign Up for Slack and Create a Slack Team](#)
- [Step 3: Create a Slack Application](#)
- [Step 4: Integrate the Slack Application with the Amazon Lex Bot](#)
- [Step 5: Complete Slack Integration](#)
- [Step 6: Test the Integration](#)

Step 1: Create an Amazon Lex Bot

If you don't already have an Amazon Lex bot, create and deploy one. In this topic, we assume that you are using the bot that you created in Getting Started Exercise 1. However, you can use any of the example bots provided in this guide. For Getting Started Exercise 1, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#)

1. Create an Amazon Lex bot. For instructions, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#).
2. Deploy the bot and create an alias. For instructions, see [Exercise 3: Publish a Version and Create an Alias](#).

Next Step

[Step 2: Sign Up for Slack and Create a Slack Team](#)

Step 2: Sign Up for Slack and Create a Slack Team

Sign up for a Slack account and create a Slack team. For instructions, see [Using Slack](#). In the next section, you create a Slack application, which any Slack team can install.

Next Step

[Step 3: Create a Slack Application](#)

Step 3: Create a Slack Application

In this section, you do the following:

1. Create a Slack application on the Slack API Console
2. Configure the application to add interactive messaging to your bot:

At the end of this section, you get application credentials (Client Id, Client Secret, and Verification Token). In the next section, you use this information to configure bot channel association in the Amazon Lex console.

1. Sign in to the Slack API Console at <http://api.slack.com>.
2. Create an application.

After you have successfully created the application, Slack displays the **Basic Information** page for the application.

3. Configure the application features as follows:
 - In the left menu, choose **Interactivity & Shortcuts**.
 - Choose the toggle to turn interactive components on.
 - In the **Request URL** box, specify any valid URL. For example, you can use **https://slack.com**.

Note

For now, enter any valid URL to get the verification token that you need in the next step. You will update this URL after you add the bot channel association in the Amazon Lex console.

- Choose **Save Changes**.
4. In the left menu, in **Settings**, choose **Basic Information**. Record the following application credentials:
 - Client ID
 - Client Secret
 - Verification Token

Next Step

[Step 4: Integrate the Slack Application with the Amazon Lex Bot](#)

Step 4: Integrate the Slack Application with the Amazon Lex Bot

Now that you have Slack application credentials, you can integrate the application with your Amazon Lex bot. To associate the Slack application with your bot, add a bot channel association in Amazon Lex.

In the Amazon Lex console, activate a bot channel association to associate the bot with your Slack application. When the bot channel association is activated, Amazon Lex returns two URLs (**Postback URL** and **OAuth URL**). Record these URLs because you need them later.

To integrate the Slack application with your Amazon Lex bot

1. Sign in to the AWS Management Console, and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose the Amazon Lex bot that you created in Step 1.
3. Choose the **Channels** tab.
4. In the left menu, choose **Slack**.
5. On the **Slack** page, provide the following:
 - Type a name. For example, BotSlackIntegration.
 - Choose "aws/lex" from the **KMS key** drop-down.
 - For **Alias**, choose the bot alias.
 - Type the **Client Id**, **Client secret**, and **Verification Token**, which you recorded in the preceding step. These are the credentials of the Slack application.

Slack

Fill in the form below and click activate to get a callback URL to use with Slack. You can generate multiple callback URLs. [Learn more](#) on steps to integrate with Slack.

Channel Name*	BotSlackAssociation	?
Channel Description	Channel for Slack	?
IAM Role	AWSServiceRoleForLexChannels Automatically created on your behalf	?
KMS Key	aws/lex	?
Alias*	BETA	?
Client Id*	Client Id	?
Client Secret*	Client Secret	?
Verification Token*	Verification Token	?
Success Page URL	Success Page URL	?

* Required Field

Callback URLs

Fill in the form above and click activate to get a callback URL. You can generate multiple callback URLs.

6. Choose **Activate**.

The console creates the bot channel association and returns two URLs (Postback URL and OAuth URL). Record them. In the next section, you update your Slack application configuration to use these endpoints as follows:

- The Postback URL is the Amazon Lex bot's endpoint that listens to Slack events. You use this URL:
 - As the request URL in the **Event Subscriptions** feature of the Slack application.
 - To replace the placeholder value for the request URL in the **Interactive Messages** feature of the Slack application.

- The OAuth URL is your Amazon Lex bot's endpoint for an OAuth handshake with Slack.

Next Step

[Step 5: Complete Slack Integration](#)

Step 5: Complete Slack Integration

In this section, use the Slack API console to complete integration of the Slack application.

1. Sign in to the Slack API console at <http://api.slack.com>. Choose the app that you created in [Step 3: Create a Slack Application](#).
2. Update the **OAuth & Permissions** feature as follows:
 - a. In the left menu, choose **OAuth & Permissions**.
 - b. In the **Redirect URLs** section, add the OAuth URL that Amazon Lex provided in the preceding step. Choose **Add a new Redirect URL**, and then choose **Save URLs**.
 - c. In the **Bot Token Scopes** section, add two permissions with the **Add an OAuth Scope** button. Filter the list with the following text:
 - **chat:write**
 - **team:read**
3. Update the **Interactivity & Shortcuts** feature by updating the **Request URL** value to the Postback URL that Amazon Lex provided in the preceding step. Enter the postback URL that you saved in step 4, and then choose **Save Changes**.
4. Subscribe to the **Event Subscriptions** feature as follows:
 - Enable events by choosing the **On** option.
 - Set the **Request URL** value to the Postback URL that Amazon Lex provided in the preceding step.
 - In the **Subscribe to Bot Events** section, subscribe to the message .im bot event to enable direct messaging between the end user and the Slack bot.
 - Save the changes.
5. Enable sending messages from the messages tab as follows:
 - From the left menu, choose **App Home**.

- In the **Show Tabs** section, choose **Allow users to send Slash commands and messages from the messages tab**.

Next Step

[Step 6: Test the Integration](#)

Step 6: Test the Integration

Now use a browser window to test the integration of Slack with your Amazon Lex bot.

1. Choose **Manage Distribution** under **Settings**. Choose **Add to Slack** to install the application. Authorize the bot to respond to messages.
2. You are redirected to your Slack team. In the left menu, in the **Direct Messages** section, choose your bot. If you don't see your bot, choose the plus icon (+) next to **Direct Messages** to search for it.
3. Engage in a chat with your Slack application, which is linked to the Amazon Lex bot. Your bot now responds to messages.

If you created the bot using Getting Started Exercise 1, you can use the example conversations provided in that exercise. For more information, see [Step 4: Add the Lambda Function as Code Hook \(Console\)](#).

Integrating an Amazon Lex Bot with Twilio Programmable SMS

This exercise provides instructions for integrating an Amazon Lex bot with the Twilio simple messaging service (SMS). You perform the following steps:

1. Create an Amazon Lex bot
2. Integrate Twilio programmable SMS with your bot Amazon Lex
3. Engage in an interaction with the Amazon Lex bot by testing the setup using the SMS service on your mobile phone
4. Test the integration

Topics

- [Step 1: Create an Amazon Lex Bot](#)

- [Step 2: Create a Twilio SMS Account](#)
- [Step 3: Integrate the Twilio Messaging Service Endpoint with the Amazon Lex Bot](#)
- [Step 4: Test the Integration](#)

Step 1: Create an Amazon Lex Bot

If you don't already have an Amazon Lex bot, create and deploy one. In this topic, we assume that you are using the bot that you created in Getting Started Exercise 1. However, you can use any of the example bots provided in this guide. For Getting Started Exercise 1, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#).

1. Create an Amazon Lex bot. For instructions, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#).
2. Deploy the bot and create an alias. For instructions, see [Exercise 3: Publish a Version and Create an Alias](#).

Step 2: Create a Twilio SMS Account

Sign up for a Twilio account and record the following account information:

- ACCOUNT SID
- AUTH TOKEN

For sign-up instructions, see <https://www.twilio.com/console>.

Step 3: Integrate the Twilio Messaging Service Endpoint with the Amazon Lex Bot

To integrate Twilio with your Amazon Lex bot

1. To associate the Amazon Lex bot with your Twilio programmable SMS endpoint, activate bot channel association in the Amazon Lex console. When the bot channel association has been activated, Amazon Lex returns a callback URL. Record this callback URL because you need it later.
 - a. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
 - b. Choose the Amazon Lex bot that you created in Step 1.

- c. Choose the **Channels** tab.
- d. In the **Chatbots** section, choose **Twilio SMS**.
- e. On the **Twilio SMS** page, provide the following information:
 - Type a name. For example, BotTwilioAssociation.
 - Choose "aws/lex" from **KMS key**.
 - For **Alias**, choose the bot alias.
 - For **Authentication Token**, type the AUTH TOKEN for your Twilio account.
 - For **Account SID**, type the ACCOUNT SID for your Twilio account.

The screenshot shows the Amazon Lex console interface for configuring a Twilio SMS channel. The 'Channels' tab is selected, and the 'Twilio SMS' configuration form is displayed. The form contains the following fields and values:

- Name:** BotTwilioAssociation
- Description:** Channel for Twilio
- IAM Role:** AWSRoleForLexChannels (Automatically created on your behalf)
- KMS key:** aws/lex
- Alias:** Beta
- Authentication Token:** Authentication Token
- Account SID:** Account SID

An **Activate** button is located below the form fields. At the bottom of the page, there is a **Test Bot** button and a scroll-to-top arrow.

- f. Choose **Activate**.

The console creates the bot channel association and returns a callback URL. Record this URL.

2. On the Twilio console, connect the Twilio SMS endpoint to the Amazon Lex bot.
 - a. Sign in to the Twilio console at <https://www.twilio.com/console>.

- b. If you don't have a Twilio SMS endpoint, create it.
- c. Update the **Inbound Settings** configuration of the messaging service by setting the **REQUEST URL** value to the callback URL that Amazon Lex provided in the preceding step.

Step 4: Test the Integration

Use your mobile phone to test the integration between Twilio SMS and your bot.

To test integration

1. Sign in to the Twilio console at <https://www.twilio.com/console> and do the following:
 - a. Verify that you have a Twilio number associated with the messaging service under **Manage Numbers**.

You send messages to this number and engage in SMS interaction with the Amazon Lex bot from your mobile phone.

- b. Verify that your mobile phone is listed as **Verified Caller ID**.

If it isn't, follow instructions on the Twilio console to enable the mobile phone that you plan to use for testing.

Now you can use your mobile phone to send messages to the Twilio SMS endpoint, which is mapped to the Amazon Lex bot.

2. Using your mobile phone, send messages to the Twilio number.

The Amazon Lex bot responds. If you created the bot using Getting Started Exercise 1, you can use the example conversations provided in that exercise. For more information, see [Step 4: Add the Lambda Function as Code Hook \(Console\)](#).

Deploying an Amazon Lex Bot in Mobile Applications

Using AWS Amplify, you can integrate your Amazon Lex bots with mobile or web applications. For more information, see [Interactions – Getting started](#) in the *AWS Amplify Docs*.

Importing and Exporting Amazon Lex Bots, Intents, and Slot Types

You can import or export a bot, intent, or slot type. For example, if you want to share a bot with a colleague in a different AWS account, you can export it, then send it to her. If you want to add multiple utterances to a bot, you can export it, add the utterances, then import it back into your account.

You can *export* bots, intents, and slot types in either Amazon Lex (to share or modify them) or an Alexa skill format. You can *import* only in Amazon Lex format.

When you export a resource, you have to export it in a format that is compatible with the service that you are exporting to, Amazon Lex or the Alexa Skills Kit. If you export a bot in Amazon Lex format, you can reimport it into your account, or an Amazon Lex user in another account can import it into his account. You can also export a bot in a format compatible with an Alexa skill. Then you can import the bot using the Alexa Skills Kit to make your bot available with Alexa. For more information, see [Exporting to an Alexa Skill](#).

When you export a bot, intent or slot type, its resources are written to a JSON file. To export a bot, intent, or slot type, you can use either the Amazon Lex console or the [GetExport](#) operation. Import a bot, intent or slot type using the [StartImport](#).

Topics

- [Exporting and Importing in Amazon Lex Format](#)
- [Exporting to an Alexa Skill](#)

Exporting and Importing in Amazon Lex Format

To export bots, intents, and slot types, from Amazon Lex with the intention of reimporting into Amazon Lex, you use create a JSON file in Amazon Lex format. You can edit your resources in this file and import it back into Amazon Lex. For example, you can add utterances to an intent and then import the changed intent back into your account. You can also use the JSON format to share a resource. For example, you can export a bot from one AWS Region and then import it into another Region. Or you can send the JSON file to a colleague to share a bot.

Topics

- [Exporting in Amazon Lex Format](#)
- [Importing in Amazon Lex Format](#)
- [JSON Format for Importing and Exporting](#)

Exporting in Amazon Lex Format

Export your Amazon Lex bots, intents, and slot types to a format that you can import to an AWS account. You can export the following resources:

- A bot, including all of the intents and custom slot types used by the bot
- An intent, including all of the custom slot types used by the intent
- A custom slot type, including all of values for the slot type

You can export only a numbered version of a resource. You can't export a resource's `$LATEST` version.

Exporting is an asynchronous process. When the export is complete, you get an Amazon S3 presigned URL. The URL provides the location of a .zip archive that contains the exported resource in JSON format.

You use either the console or the [GetExport](#) operation to export bots, intents, and custom slot types.

The process for exporting, a bot, an intent, or a slot type is the same. In the following procedures, substitute intent or slot type for bot.

Exporting a Bot

To export a bot

1. Sign in to the AWS Management Console, and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose **Bots**, then choose the bot to export.
3. On the **Actions** menu, choose **Export**.
4. In the **Export Bot** dialog, choose the version of the bot to export. For **Platform**, choose **Amazon Lex**.

5. Choose **Export**.
6. Download and save the .zip archive.

Amazon Lex exports the bot to a JSON file that is contained in the .zip archive. To update the bot, modify the JSON text, then import it back into Amazon Lex.

Next step

[Importing in Amazon Lex Format](#)

Importing in Amazon Lex Format

After you have exported a resource to a JSON file in the Amazon Lex format, you can import the JSON file containing the resource into one or more AWS accounts. For example, you can export a bot, and then import it into another AWS Region. Or you can send the bot to a colleague so that she can import it into her account.

When you import a bot, intent, or slot type, you must decide whether you want to overwrite the \$LATEST version of a resource, such as an intent or a slot type, during import, or if you want the import to fail if you want to preserve the resource that is in your account. For example, if you are uploading an edited version of a resource to your account, you would choose to overwrite the \$LATEST version. If you are uploading a resource sent to you by a colleague, you can choose to have the import fail if there are resource conflicts so that your own resources aren't replaced.

When importing a resource, the permissions assigned to the user making the import request apply. The user must have permissions for all of the resources in the account that the import affects. The user must also have permission for the [GetBot](#), [PutBot](#), [GetIntent](#), [PutIntent](#), [GetSlotType](#), [PutSlotType](#) operations. For more information about permissions, see [How Amazon Lex works with IAM](#).

The import reports errors that occur during processing. Some errors are reported before the import begins, others are reported during the import process. For example, if the account that is importing an intent doesn't have permission to call a Lambda function that the intent uses, the import fails before changes are made to the slot types or intents. If an import fails during the import process, the \$LATEST version of any intent or slot type imported before the process failed is modified. You can't roll back changes made to the \$LATEST version.

When you import a resource, all dependent resources are imported to the \$LATEST version of the resource and then given a numbered version. For example, if a bot uses an intent, the intent is

given a numbered version. If an intent uses a custom slot type, the slot type is given a numbered version.

A resource is imported only once. For example, if the bot contains an `OrderPizza` intent and an `OrderDrink` intent that both rely on the custom slot type `Size`, the `Size` slot type is imported once and used for both intents.

Note

If you exported your bot with the `enableModelImprovements` parameter set to `false`, you must open the .zip file containing the bot definition and change the `enableModelImprovements` parameter to `true` in the following Regions:

- Asia Pacific (Singapore) (ap-southeast-1)
- Asia Pacific (Tokyo) (ap-northeast-1)
- EU (Frankfurt) (eu-central-1)
- EU (London) (eu-west-2)

The process for importing a bot, an intent, or a custom slot type is the same. In the following procedures, substitute intent or slot type, as appropriate.

Importing a Bot

To import a bot

1. Sign in to the AWS Management Console, and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose **Bots**, then choose the bot to import. To import a new bot, skip this step.
3. For **Actions**, choose **Import**.
4. For **Import Bot**, choose the .zip archive that contains the JSON file that contains the bot to import. If you want to see merge conflicts before merging, choose **Notify me of merge conflicts**. If you turn off conflict checking, the \$LATEST version of all of the resources used by the bot are overwritten.
5. Choose **Import**. If you have chosen to be notified of merge conflicts and there are conflicts, a dialog appears that lists them. To overwrite the \$LATEST version of all conflicting resources, choose **Overwrite and continue**. To stop the import, choose **Cancel**.

You can now test the bot in your account.

JSON Format for Importing and Exporting

The following examples show the JSON structure for exporting and importing slot types, intents, and bots in Amazon Lex format.

Slot Type structure

The following is the JSON structure for custom slot types. Use this structure when you import or export slot types, and when you export intents that depend on custom slot types.

```
{
  "metadata": {
    "schemaVersion": "1.0",
    "importType": "LEX",
    "importFormat": "JSON"
  },
  "resource": {
    "name": "slot type name",
    "version": "version number",
    "enumerationValues": [
      {
        "value": "enumeration value",
        "synonyms": []
      },
      {
        "value": "enumeration value",
        "synonyms": []
      }
    ],
    "valueSelectionStrategy": "ORIGINAL_VALUE or TOP_RESOLUTION"
  }
}
```

Intent structure

The following is the JSON structure for intents. Use this structure when you import or export intents and bots that depend on an intent.

```
{
  "metadata": {
    "schemaVersion": "1.0",
```

```
"importType": "LEX",
"importFormat": "JSON"
},
"resource": {
  "description": "intent description",
  "rejectionStatement": {
    "messages": [
      {
        "contentType": "PlainText or SSML or CustomPayload",
        "content": "string"
      }
    ]
  },
  "name": "intent name",
  "version": "version number",
  "fulfillmentActivity": {
    "type": "ReturnIntent or CodeHook"
  },
  "sampleUtterances": [
    "string",
    "string"
  ],
  "slots": [
    {
      "name": "slot name",
      "description": "slot description",
      "slotConstraint": "Required or Optional",
      "slotType": "slot type",
      "valueElicitationPrompt": {
        "messages": [
          {
            "contentType": "PlainText or SSML or CustomPayload",
            "content": "string"
          }
        ]
      },
      "maxAttempts": value
    },
    {
      "priority": value,
      "sampleUtterances": []
    }
  ],
  "confirmationPrompt": {
    "messages": [
      {
```

```

        "contentType": "PlainText or SSML or CustomPayload",
        "content": "string"
    },
    {
        "contentType": "PlainText or SSML or CustomPayload",
        "content": "string"
    }
],
"maxAttempts": value
},
"slotTypes": [
    List of slot type JSON structures.
    For more information, see Slot Type structure.
]
}
}

```

Bot structure

The following is the JSON structure for bots. Use this structure when you import or export bots.

```

{
  "metadata": {
    "schemaVersion": "1.0",
    "importType": "LEX",
    "importFormat": "JSON"
  },
  "resource": {
    "name": "bot name",
    "version": "version number",
    "nluIntentConfidenceThreshold": 0.00-1.00,
    "enableModelImprovements": true | false,
    "intents": [
      List of intent JSON structures.
      For more information, see Intent structure.
    ],
    "slotTypes": [
      List of slot type JSON structures.
      For more information, see Slot Type structure.
    ],
    "voiceId": "output voice ID",
    "childDirected": boolean,
    "locale": "en-US",

```

```
"idleSessionTTLInSeconds": timeout,
"description": "bot description",
"clarificationPrompt": {
  "messages": [
    {
      "contentType": "PlainText or SSML or CustomPayload",
      "content": "string"
    }
  ],
  "maxAttempts": value
},
"abortStatement": {
  "messages": [
    {
      "contentType": "PlainText or SSML or CustomPayload",
      "content": "string"
    }
  ]
}
}
```

Exporting to an Alexa Skill

You can export your bot schema in a format compatible with an Alexa skill. After you export the bot to a JSON file, you upload it to Alexa using the skill builder.

To export a bot and its schema (interaction model)

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. Choose the bot that you want to export.
3. For **Actions**, choose **Export**.
4. Choose the version of the bot that you want to export. For the format, choose **Alexa Skills Kit**, then choose **Export**.
5. If a download dialog box appears, choose a location to save the file, then choose **Save**.

The downloaded file is a .zip archive containing one file with the name of the exported bot. It contains the information necessary to import the bot as an Alexa skill.

 Note

Amazon Lex and the Alexa Skills Kit differ in the following ways:

- Session attributes, denoted by square brackets ([]), are not supported by the Alexa Skills Kit. You need to update prompts that use session attributes.
- Punctuation marks are not supported by the Alexa Skills Kit. You need to update utterances that use punctuation.

To upload the bot to an Alexa Skill

1. Log in to the developer portal at <https://developer.amazon.com/>.
2. On the **Alexa Skills** page, choose **Create Skill**.
3. On the **Create a new skill** page, enter a skill name and the default language for the skill. Make sure that **Custom** is selected for the skill model, and then choose **Create skill**.
4. Make sure that **Start from scratch** is selected and then choose **Choose**.
5. From the left menu, choose **JSON Editor**. Drag the JSON file that you exported from Amazon Lex to the JSON editor.
6. Choose **Save Model** to save your interaction model.

After uploading the schema into the Alexa skill, make changes necessary for running the skill with Alexa. For more information about creating an Alexa skill, see [Use the Skill Builder \(Beta\)](#) in the *Alexa Skills Kit*.

Additional Examples: Creating Amazon Lex Bots

The following sections provide additional Amazon Lex exercises with step-by-step instructions.

Topics


- [Schedule Appointment](#)
- [Book Trip](#)
- [Using a Response Card](#)
- [Updating Utterances](#)
- [Integrating with a Web site](#)
- [Call Center Agent Assistant](#)

Schedule Appointment


The example bot in this exercise schedules appointments for a dentist's office. The example also illustrates using response cards to obtain user input with buttons. Specifically, the example illustrates generating response cards dynamically at runtime.

You can configure response cards at build time (also referred to as static response cards) or generate them dynamically in an AWS Lambda function. In this example, the bot uses the following response cards:

- A response card that lists buttons for appointment type. See the following image for an example:



- A response card that lists buttons for appointment date. See the following image for an example:




When would you like to schedule your root canal?

Specify Date
When would you like to schedule your root canal?

2-15 (Wed)
2-16 (Thu)
2-17 (Fri)

This image shows a response card with a light gray header containing the question "When would you like to schedule your root canal?". Below the header is a white box with a rounded top-left corner. Inside this box, the text "Specify Date" is in bold, followed by the same question. Below the text is a table with three rows, each containing a date and day of the week: "2-15 (Wed)", "2-16 (Thu)", and "2-17 (Fri)".

- A response card that lists buttons to confirm a suggested appointment time. See the following image for an example:



What time on 2017-02-15 works for you? 4:00 p.m. is our only availability, does that work for you?

Confirm Appointment
Is 4:00 p.m. on 2017-02-15 okay?

yes
no

This image shows a response card with a light gray header containing the question "What time on 2017-02-15 works for you? 4:00 p.m. is our only availability, does that work for you?". Below the header is a white box with a rounded top-left corner. Inside this box, the text "Confirm Appointment" is in bold, followed by the question "Is 4:00 p.m. on 2017-02-15 okay?". Below the text is a table with two rows, each containing a response option: "yes" and "no".

The available appointment dates and times vary, which requires you to generate response cards at runtime. You use an AWS Lambda function to generate these response cards dynamically. The Lambda function returns response cards in its response to Amazon Lex. Amazon Lex includes the response card in its response to the client.

If a client (for example, Facebook Messenger) supports response cards, the user can either choose from the list of buttons or type the response. Otherwise, the user simply types the response.

In addition to the button shown in the preceding example, you can also include images, attachments, and other useful information to display on response cards. For information about response cards, see [Response Cards](#).

In this exercise, you do the following:

- Create and test a bot (using the ScheduleAppointment blueprint). For this exercise, you use a bot blueprint to quickly set up and test the bot. For a list of available blueprints, see [Amazon Lex and AWS Lambda Blueprints](#). This bot is preconfigured with one intent (MakeAppointment).
- Create and test a Lambda function (using the lex-make-appointment-python blueprint provided by Lambda). You configure the MakeAppointment intent to use this Lambda function as a code hook to perform initialization, validation, and fulfillment tasks.

 Note

The provided example Lambda function showcases a dynamic conversation based on the mocked-up availability of a dentist appointment. In a real application, you might use a real calendar to set an appointment.

- Update the MakeAppointment intent configuration to use the Lambda function as a code hook. Then, test the end-to-end experience.
- Publish the schedule appointment bot to Facebook Messenger so you can see the response cards in action (the client in the Amazon Lex console currently does not support response cards).

The following sections provide summary information about the blueprints you use in this exercise.

Topics

- [Overview of the Bot Blueprint \(ScheduleAppointment\)](#)
- [Overview of the Lambda Function Blueprint \(lex-make-appointment-python\)](#)
- [Step 1: Create an Amazon Lex Bot](#)
- [Step 2: Create a Lambda Function](#)
- [Step 3: Update the Intent: Configure a Code Hook](#)
- [Step 4: Deploy the Bot on the Facebook Messenger Platform](#)
- [Details of Information Flow](#)

Overview of the Bot Blueprint (ScheduleAppointment)

The ScheduleAppointment blueprint that you use to create a bot for this exercise is preconfigured with the following:

- **Slot types** – One custom slot type called AppointmentTypeValue, with the enumeration values root canal, cleaning, and whitening.
- **Intent** – One intent (MakeAppointment), which is preconfigured as follows:
 - **Slots** – The intent is configured with the following slots:
 - Slot AppointmentType, of the AppointmentTypes custom type.
 - Slot Date, of the AMAZON.DATE built-in type.
 - Slot Time, of the AMAZON.TIME built-in type.
 - **Utterances** – The intent is preconfigured with the following utterances:
 - "I would like to book an appointment"
 - "Book an appointment"
 - "Book a {AppointmentType}"

If the user utters any of these, Amazon Lex determines that MakeAppointment is the intent, and then uses the prompts to elicit slot data.

- **Prompts** – The intent is preconfigured with the following prompts:
 - Prompt for the AppointmentType slot – "What type of appointment would you like to schedule?"
 - Prompt for the Date slot – "When should I schedule your {AppointmentType}?"
 - Prompt for the Time slot – "At what time do you want to schedule the {AppointmentType}?" and
"At what time on {Date}?"
 - Confirmation prompt – "{Time} is available, should I go ahead and book your appointment?"
 - Cancel message– "Okay, I will not schedule an appointment."

Overview of the Lambda Function Blueprint (lex-make-appointment-python)

The Lambda function blueprint (lex-make-appointment-python) is a code hook for bots that you create using the ScheduleAppointment bot blueprint.

This Lambda function blueprint code can perform both initialization/validation and fulfillment tasks.

- The Lambda function code showcases a dynamic conversation that is based on example availability for a dentist appointment (in real applications, you might use a calendar). For the day or date that the user specifies, the code is configured as follows:
 - If there are no appointments available, the Lambda function returns a response directing Amazon Lex to prompt the user for another day or date (by setting the `dialogAction` type to `ElicitSlot`). For more information, see [Response Format](#).
 - If there is only one appointment available on the specified day or date, the Lambda function suggests the available time in the response and directs Amazon Lex to obtain user confirmation by setting the `dialogAction` in the response to `ConfirmIntent`. This illustrates how you can improve the user experience by proactively suggesting the available time for an appointment.
 - If there are multiple appointments available, the Lambda function returns a list of available times in the response to Amazon Lex. Amazon Lex returns a response to the client with the message from the Lambda function.
- As the fulfillment code hook, the Lambda function returns a summary message indicating that an appointment is scheduled (that is, the intent is fulfilled).

Note

In this example, we show how to use response cards. The Lambda function constructs and returns a response card to Amazon Lex. The response card lists available days and times as buttons to choose from. When testing the bot using the client provided by the Amazon Lex console, you cannot see the response card. To see it, you must integrate the bot with a messaging platform, such as Facebook Messenger. For instructions, see [Integrating an Amazon Lex Bot with Facebook Messenger](#). For more information about response cards, see [Managing Messages](#).

When Amazon Lex invokes the Lambda function, it passes event data as input. One of the event fields is `invocationSource`, which the Lambda function uses to choose between an input validation and fulfillment activity. For more information, see [Input Event Format](#).

Next Step

[Step 1: Create an Amazon Lex Bot](#)

Step 1: Create an Amazon Lex Bot

In this section, you create an Amazon Lex bot using the `ScheduleAppointment` blueprint, which is provided in the Amazon Lex console.

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. On the **Bots** page, choose **Create**.
3. On the **Create your Lex bot** page, do the following:
 - Choose the **ScheduleAppointment** blueprint.
 - Leave the default bot name (`ScheduleAppointment`).
4. Choose **Create**.


This step saves and builds the bot. The console sends the following requests to Amazon Lex during the build process:

- Create a new version of the slot types (from the `$LATEST` version). For information about slot types defined in this bot blueprint, see [Overview of the Bot Blueprint \(ScheduleAppointment\)](#).
- Create a version of the `MakeAppointment` intent (from the `$LATEST` version). In some cases, the console sends a request for the update API operation before creating a new version.
- Update the `$LATEST` version of the bot.

At this time, Amazon Lex builds a machine learning model for the bot. When you test the bot in the console, the console uses the runtime API to send user input back to Amazon Lex. Amazon Lex then uses the machine learning model to interpret the user input.

5. The console shows the `ScheduleAppointment` bot. On the **Editor** tab, review the preconfigured intent (`MakeAppointment`) details.

6. Test the bot in the test window. Use the following screen shot to engage in a test conversation with your bot:



Note the following:

- From the initial user input ("Book an appointment"), the bot infers the intent (MakeAppointment).
- The bot then uses the configured prompts to get slot data from the user.
- The bot blueprint has the MakeAppointment intent configured with the following confirmation prompt:

```
{Time} is available, should I go ahead and book your appointment?
```

After the user provides all of the slot data, Amazon Lex returns a response to the client with a confirmation prompt as the message. The client displays the message for the user:

16:00 is available, should I go ahead and book your appointment?

Notice that the bot accepts any appointment date and time values because you don't have any code to initialize or validate the user data. In the next section, you add a Lambda function to do this.

Next Step

[Step 2: Create a Lambda Function](#)

Step 2: Create a Lambda Function

In this section, you create a Lambda function using a blueprint (lex-make-appointment-python) that is provided in the Lambda console. You also test the Lambda function by invoking it using sample Amazon Lex event data that is provided by the console.

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. Choose **Create a Lambda function**.
3. For **Select blueprint**, type **lex** to find the blueprint, and then choose the **lex-make-appointment-python** blueprint.
4. Configure the Lambda function as follows.
 - Type the Lambda function name (MakeAppointmentCodeHook).
 - For the role, choose **Create a new role from template(s)**, and then type a role name.
 - Leave other default values.
5. Choose **Create Function**.
6. If you are using a locale other than English (US) (en-US), update the intent names as described in [Updating a Blueprint for a Specific Locale](#).
7. Test the Lambda function.
 - a. Choose **Actions**, and then choose **Configure test event**.

- b. From the **Sample event template** list, choose **Lex-Make Appointment (preview)**. This sample event uses the Amazon Lex request/response model, with values set to match a request from your Amazon Lex bot. For information about the Amazon Lex request/response model, see [Using Lambda Functions](#).
- c. Choose **Save and test**.
- d. Verify that the Lambda function ran successfully. The response in this case matches the Amazon Lex response model.

Next Step

[Step 3: Update the Intent: Configure a Code Hook](#)

Step 3: Update the Intent: Configure a Code Hook

In this section, you update the configuration of the MakeAppointment intent to use the Lambda function as a code hook for the validation and fulfillment activities.

1. In the Amazon Lex console, select the ScheduleAppointment bot. The console shows the **MakeAppointment** intent. Modify the intent configuration as follows.

Note

You can update only the \$LATEST versions of any of the Amazon Lex resources, including the intents. Make sure that the intent version is set to \$LATEST. You have not published a version of your bot yet, so it should still be the \$LATEST version in the console.

- a. In the **Options** section, choose **Initialization and validation code hook**, and then choose the Lambda function from the list.
 - b. In the **Fulfillment** section, choose **AWS Lambda function**, and then choose the Lambda function from the list.
 - c. Choose **Goodbye message**, and type a message.
2. Choose **Save**, and then choose **Build**.
 3. Test the bot, as in the following image:

The screenshot shows the 'Test Bot' interface in the Amazon Lex console. At the top, it says 'Build: Latest | Status: READY'. The conversation starts with the user input 'Book appointment'. The bot responds with 'What type of appointment would you like to schedule?'. The user inputs 'root canal'. The bot asks 'When would you like to schedule your root canal?'. The user inputs 'Tuesday'. The bot responds with 'We do not have any availability on that date, is there another day which works for you?'. The user inputs 'Wednesday'. The bot asks 'What time on 2017-01-18 works for you? 4:00 p.m. is our only availability, does that work for you?'. The user inputs 'yes'. The bot responds with 'Okay, I have booked your appointment. We will see you at 4:00 p.m. on 2017-01-18'. At the bottom, there is a 'Clear' button and a text input field with the placeholder 'Type to your bot...'.

Next Step

[Step 4: Deploy the Bot on the Facebook Messenger Platform](#)

Step 4: Deploy the Bot on the Facebook Messenger Platform

In the preceding section, you tested the ScheduleAppointment bot using the client in the Amazon Lex console. Currently, the Amazon Lex console does not support response cards. To test the dynamically generated response cards that the bot supports, deploy the bot on the Facebook Messenger platform and test it.

For instructions, see [Integrating an Amazon Lex Bot with Facebook Messenger](#).

Next Step

[Details of Information Flow](#)

Details of Information Flow

The ScheduleAppointment bot blueprint primarily showcases the use of dynamically generated response cards. The Lambda function in this exercise includes response cards in its response to Amazon Lex. Amazon Lex includes the response cards in its reply to the client. This section explains both the following:

- Data flow between client and Amazon Lex.

The section assumes client sends requests to Amazon Lex using the PostText runtime API and shows request/response details accordingly. For more information about the PostText runtime API, see [PostText](#).

Note

For an example of information flow between client and Amazon Lex in which client uses the PostContent API, see [Step 2a \(Optional\): Review the Details of the Spoken Information Flow \(Console\)](#).

- Data flow between Amazon Lex and the Lambda function. For more information, see [Lambda Function Input Event and Response Format](#).

Note

The example assumes that you are using the Facebook Messenger client, which does not pass session attributes in the request to Amazon Lex. Accordingly, the example requests shown in this section show empty sessionAttributes. If you test the bot using the client provided in the Amazon Lex console, the client includes the session attributes.

This section describes what happens after each user input.

1. User: Types **Book an appointment**.

- a. The client (console) sends the following [PostContent](#) request to Amazon Lex:

```
POST /bot/ScheduleAppointment/alias/$LATEST/
user/bijt6rovckwecnzsbthrr1d7lv3ja3n/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText":"book appointment",
  "sessionAttributes":{}
}
```

Both the request URI and the body provide information to Amazon Lex:

- Request URI – Provides the bot name (*ScheduleAppointment*), the bot alias (*\$LATEST*), and the user name ID. The trailing text indicates that it is a PostText (not PostContent) API request.
 - Request body – Includes the user input (*inputText*) and empty *sessionAttributes*.
- b. From the *inputText*, Amazon Lex detects the intent (*MakeAppointment*). The service invokes the Lambda function, which is configured as a code hook, to perform initialization and validation by passing the following event. For details, see [Input Event Format](#).

```
{
  "currentIntent": {
    "slots": {
      "AppointmentType": null,
      "Date": null,
      "Time": null
    },
    "name": "MakeAppointment",
    "confirmationStatus": "None"
  },
  "bot": {
    "alias": null,
    "version": "$LATEST",
    "name": "ScheduleAppointment"
  },
}
```

```
"userId": "bijt6rovckwecnzesbthrr1d7lv3ja3n",
"invocationSource": "DialogCodeHook",
"outputDialogMode": "Text",
"messageVersion": "1.0",
"sessionAttributes": {}
}
```

In addition to the information sent by the client, Amazon Lex also includes the following data:


- **currentIntent** – Provides current intent information.
 - **invocationSource** – Indicates the purpose of the Lambda function invocation. In this case, the purpose is to perform user data initialization and validation. (Amazon Lex knows that the user has not provided all of the slot data to fulfill the intent yet.)
 - **messageVersion** – Currently Amazon Lex supports only the 1.0 version.
- c. At this time, all of the slot values are null (there is nothing to validate). The Lambda function returns the following response to Amazon Lex, directing the service to elicit information for the **AppointmentType** slot. For information about the response format, see [Response Format](#).

```
{
  "dialogAction": {
    "slotToElicit": "AppointmentType",
    "intentName": "MakeAppointment",
    "responseCard": {
      "genericAttachments": [
        {
          "buttons": [
            {
              "text": "cleaning (30 min)",
              "value": "cleaning"
            },
            {
              "text": "root canal (60 min)",
              "value": "root canal"
            },
            {
              "text": "whitening (30 min)",
              "value": "whitening"
            }
          ]
        }
      ]
    }
  }
}
```


```
        "subTitle": "What type of appointment would you like to  
schedule?",  
        "title": "Specify Appointment Type"  
    },  
    ],  
    "version": 1,  
    "contentType": "application/vnd.amazonaws.card.generic"  
},  
"slots": {  
    "AppointmentType": null,  
    "Date": null,  
    "Time": null  
},  
"type": "ElicitSlot",  
"message": {  
    "content": "What type of appointment would you like to schedule?",  
    "contentType": "PlainText"  
}  
},  
"sessionAttributes": {}  
}
```

The response includes the `dialogAction` and `sessionAttributes` fields. Among other things, the `dialogAction` field returns the following fields:

- **type** – By setting this field to `ElicitSlot`, the Lambda function directs Amazon Lex to elicit the value for the slot specified in the `slotToElicit` field. The Lambda function also provides a message to convey to the user.
- **responseCard** – Identifies a list of possible values for the `AppointmentType` slot. A client that supports response cards (for example, the Facebook Messenger) displays a response card to allow the user to choose an appointment type, as in the following image:



- d. As indicated by the `dialogAction.type` in the response from the Lambda function, Amazon Lex sends the following response back to the client:



The client reads the response, and then displays the message: "What type of appointment would you like to schedule?" and the response card (if the client supports response cards).

2. User: Depending on the client, the user has two options:
 - If the response card is shown, choose **root canal (60 min)** or type **root canal**.
 - If the client does not support response cards, type **root canal**.

- a. The client sends the following PostText request to Amazon Lex (line breaks have been added for readability):

```
POST /bot/BookTrip/alias/$LATEST/user/bijt6rovckwecnzeshthrr1d7lv3ja3n/text
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"

{
  "inputText": "root canal",
  "sessionAttributes": {}
}
```

- b. Amazon Lex invokes the Lambda function for user data validation by sending the following event as a parameter:

```
{
  "currentIntent": {
    "slots": {
      "AppointmentType": "root canal",
      "Date": null,
      "Time": null
    },
    "name": "MakeAppointment",
    "confirmationStatus": "None"
  },
  "bot": {
    "alias": null,
    "version": "$LATEST",
    "name": "ScheduleAppointment"
  },
  "userId": "bijt6rovckwecnzeshthrr1d7lv3ja3n",
  "invocationSource": "DialogCodeHook",
  "outputDialogMode": "Text",
  "messageVersion": "1.0",
  "sessionAttributes": {}
}
```

In the event data, note the following:

- `invocationSource` continues to be `DialogCodeHook`. In this step, we are just validating user data.
 - Amazon Lex sets the `AppointmentType` field in the `currentIntent.slots` slot to `root canal`.
 - Amazon Lex simply passes the `sessionAttributes` field between the client and the Lambda function.
- c. The Lambda function validates the user input and returns the following response to Amazon Lex, directing the service to elicit a value for the appointment date.


```
{
  "dialogAction": {
    "slotToElicit": "Date",
    "intentName": "MakeAppointment",
    "responseCard": {
      "genericAttachments": [
        {
          "buttons": [
            {
              "text": "2-15 (Wed)",
              "value": "Wednesday, February 15, 2017"
            },
            {
              "text": "2-16 (Thu)",
              "value": "Thursday, February 16, 2017"
            },
            {
              "text": "2-17 (Fri)",
              "value": "Friday, February 17, 2017"
            },
            {
              "text": "2-20 (Mon)",
              "value": "Monday, February 20, 2017"
            },
            {
              "text": "2-21 (Tue)",
              "value": "Tuesday, February 21, 2017"
            }
          ],
          "subTitle": "When would you like to schedule your root canal?",

```

```
        "title": "Specify Date"
      },
    ],
    "version": 1,
    "contentType": "application/vnd.amazonaws.card.generic"
  },
  "slots": {
    "AppointmentType": "root canal",
    "Date": null,
    "Time": null
  },
  "type": "ElicitSlot",
  "message": {
    "content": "When would you like to schedule your root canal?",
    "contentType": "PlainText"
  }
},
"sessionAttributes": {}
}
```

Again, the response includes the `dialogAction` and `sessionAttributes` fields. Among other things, the `dialogAction` field returns the following fields:

- **type** – By setting this field to `ElicitSlot`, the Lambda function directs Amazon Lex to elicit the value for the slot specified in the `slotToElicit` field. The Lambda function also provides a message to convey to the user.
- **responseCard** – Identifies a list of possible values for the Date slot. A client that supports response cards (for example, Facebook Messenger) displays a response card that allows the user to choose an appointment date, as in the following image:



When would you like to schedule your root canal?


Specify Date
When would you like to schedule your root canal?

2-15 (Wed)
2-16 (Thu)
2-17 (Fri)

Although the Lambda function returned five dates, the client (Facebook Messenger) has a limit of three buttons for a response card. Therefore, you see only the first three values in the screen shot.

These dates are hard coded in the Lambda function. In a production application, you might use a calendar to get available dates in real time. Because the dates are dynamic, you must generate the response card dynamically in the Lambda function.

- d. Amazon Lex notices the `dialogAction.type` and returns the following response to the client that includes information from the Lambda function's response.



Headers	Cookies	Params	Response
Filter properties			
▼ JSON			
<pre>dialogState: "ElicitSlot" intentName: "MakeAppointment" message: "When would you like to schedule your root canal?" responseCard: Object contentType: "application/vnd.amazonaws.card.generic" genericAttachments: Object version: "1" sessionAttributes: Object slotToElicit: "Date" slots: Object AppointmentType: "root canal" Date: null Time: null</pre>			

The client displays the message: **When would you like to schedule your root canal?** and the response card (if the client supports response cards).

3. User: Types **Thursday**.

- a. The client sends the following PostText request to Amazon Lex (line breaks have been added for readability):

```
POST /bot/BookTrip/alias/$LATEST/user/bijt6rovckwecnzsbthrr1d7lv3ja3n/text
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"

{
  "inputText": "Thursday",
  "sessionAttributes": {}
}
```

- b. Amazon Lex invokes the Lambda function for user data validation by sending in the following event as a parameter:

```
{
  "currentIntent": {
    "slots": {
      "AppointmentType": "root canal",
      "Date": "2017-02-16",
      "Time": null
    },
    "name": "MakeAppointment",
    "confirmationStatus": "None"
  },
  "bot": {
    "alias": null,
    "version": "$LATEST",
    "name": "ScheduleAppointment"
  },
  "userId": "u3fpr9gghj02zts7y5tpq5mm4din2xqy",
  "invocationSource": "DialogCodeHook",
  "outputDialogMode": "Text",
  "messageVersion": "1.0",
  "sessionAttributes": {}
}
```

In the event data, note the following:

- `invocationSource` continues to be `DialogCodeHook`. In this step, we are just validating the user data.
 - Amazon Lex sets the `Date` field in the `currentIntent.slots` slot to `2017-02-16`.
 - Amazon Lex simply passes the `sessionAttributes` between the client and the Lambda function.
- c. The Lambda function validates the user input. This time the Lambda function determines that there are no appointments available on the specified date. It returns the following response to Amazon Lex, directing the service to again elicit a value for the appointment date.

```
{
  "dialogAction": {
    "slotToElicit": "Date",
    "intentName": "MakeAppointment",
    "responseCard": {
      "genericAttachments": [
        {
          "buttons": [
            {
              "text": "2-15 (Wed)",
              "value": "Wednesday, February 15, 2017"
            },
            {
              "text": "2-17 (Fri)",
              "value": "Friday, February 17, 2017"
            },
            {
              "text": "2-20 (Mon)",
              "value": "Monday, February 20, 2017"
            },
            {
              "text": "2-21 (Tue)",
              "value": "Tuesday, February 21, 2017"
            }
          ],
          "subTitle": "When would you like to schedule your root canal?",
          "title": "Specify Date"
        }
      ]
    }
  }
}
```


```

        }
      ],
      "version": 1,
      "contentType": "application/vnd.amazonaws.card.generic"
    },
    "slots": {
      "AppointmentType": "root canal",
      "Date": null,
      "Time": null
    },
    "type": "ElicitSlot",
    "message": {
      "content": "We do not have any availability on that date, is there another day which works for you?",
      "contentType": "PlainText"
    }
  },
  "sessionAttributes": {
    "bookingMap": "{\"2017-02-16\": []}"
  }
}

```

Again, the response includes the `dialogAction` and `sessionAttributes` fields. Among other things, the `dialogAction` returns the following fields:

- `dialogAction` field:
 - `type` – The Lambda function sets this value to `ElicitSlot` and resets the `slotToElicit` field to `Date`. The Lambda function also provides an appropriate message to convey to the user.
 - `responseCard` – Returns a list of values for the `Date` slot.
 - `sessionAttributes` - This time the Lambda function includes the `bookingMap` session attribute. Its value is the requested date of the appointment and available appointments (an empty object indicates that no appointments are available).
- d. Amazon Lex notices the `dialogAction.type` and returns the following response to the client that includes information from the Lambda function's response.



The client displays the message: **We do not have any availability on that date, is there another day which works for you?** and the response card (if the client supports response cards).

4. User: Depending on the client, the user has two options:
 - If the response card is shown, choose **2-15 (Wed)** or type **Wednesday**.
 - If the client does not support response cards, type **Wednesday**.
- a. The client sends the following PostText request to Amazon Lex:

```

POST /bot/BookTrip/alias/$LATEST/user/bijt6rovckwecnzsbthrr1d7lv3ja3n/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText": "Wednesday",
  "sessionAttributes": {
  }
}

```

Note

The Facebook Messenger client does not set any session attributes. If you want to maintain session states between requests, you must do so in the Lambda function. In a real application, you might need to maintain these session attributes in a backend database.

- b. Amazon Lex invokes the Lambda function for user data validation by sending the following event as a parameter:

```
{
  "currentIntent": {
    "slots": {
      "AppointmentType": "root canal",
      "Date": "2017-02-15",
      "Time": null
    },
    "name": "MakeAppointment",
    "confirmationStatus": "None"
  },
  "bot": {
    "alias": null,
    "version": "$LATEST",
    "name": "ScheduleAppointment"
  },
  "userId": "u3fpr9gghj02zts7y5tpq5mm4din2xqy",
  "invocationSource": "DialogCodeHook",
  "outputDialogMode": "Text",
  "messageVersion": "1.0",
  "sessionAttributes": {
  }
}
```

Amazon Lex updated `currentIntent.slots` by setting the `Date` slot to `2017-02-15`.

- c. The Lambda function validates the user input and returns the following response to Amazon Lex, directing it to elicit the value for the appointment time.

```
{
  "dialogAction": {
    "slots": {
```

```


        "AppointmentType": "root canal",
        "Date": "2017-02-15",
        "Time": "16:00"
    },
    "message": {
        "content": "What time on 2017-02-15 works for you? 4:00 p.m. is our
only availability, does that work for you?",
        "contentType": "PlainText"
    },
    "type": "ConfirmIntent",
    "intentName": "MakeAppointment",
    "responseCard": {
        "genericAttachments": [
            {
                "buttons": [
                    {
                        "text": "yes",
                        "value": "yes"
                    },
                    {
                        "text": "no",
                        "value": "no"
                    }
                ],
                "subTitle": "Is 4:00 p.m. on 2017-02-15 okay?",
                "title": "Confirm Appointment"
            }
        ],
        "version": 1,
        "contentType": "application/vnd.amazonaws.card.generic"
    }
},
"sessionAttributes": {
    "bookingMap": "{\"2017-02-15\": [\"10:00\", \"16:00\", \"16:30\"]}"
}
}

```


Again, the response includes the `dialogAction` and `sessionAttributes` fields. Among other things, the `dialogAction` returns the following fields:

- `dialogAction` field:

- **type** – The Lambda function sets this value to `ConfirmIntent`, directing Amazon Lex to obtain user confirmation of the appointment time suggested in the message.
- **responseCard** – Returns a list of yes/no values for the user to choose from. If the client supports response cards, it displays the response card, as shown in the following example:



- **sessionAttributes** - The Lambda function sets the `bookingMap` session attribute with its value set to the appointment date and available appointments on that date. In this example, these are 30-minute appointments. For a root canal that requires one hour, only 4 p.m. can be booked.
- d. As indicated in the `dialogAction.type` in the Lambda function's response, Amazon Lex returns the following response to the client:



The client displays the message: **What time on 2017-02-15 works for you? 4:00 p.m. is our only availability, does that work for you?**

5. User: Choose **yes**.

Amazon Lex invokes the Lambda function with the following event data. Because the user replied **yes**, Amazon Lex sets the `confirmationStatus` to `Confirmed`, and sets the `Time` field in `currentIntent.slots` to 4 p.m.

```

{
  "currentIntent": {
    "slots": {
      "AppointmentType": "root canal",
      "Date": "2017-02-15",
      "Time": "16:00"
    },
    "name": "MakeAppointment",
    "confirmationStatus": "Confirmed"
  },
  "bot": {
    "alias": null,
    "version": "$LATEST",

```



```
    "name": "ScheduleAppointment"
  },
  "userId": "u3fpr9gghj02zts7y5tpq5mm4din2xqy",
  "invocationSource": "FulfillmentCodeHook",
  "outputDialogMode": "Text",
  "messageVersion": "1.0",
  "sessionAttributes": {
  }
}
```

Because the `confirmationStatus` is confirmed, the Lambda function processes the intent (books a dental appointment) and returns the following response to Amazon Lex:

```
{
  "dialogAction": {
    "message": {
      "content": "Okay, I have booked your appointment.  We will see you at 4:00 p.m. on 2017-02-15",
      "contentType": "PlainText"
    },
    "type": "Close",
    "fulfillmentState": "Fulfilled"
  },
  "sessionAttributes": {
    "formattedTime": "4:00 p.m.",
    "bookingMap": "{\"2017-02-15\": [\"10:00\"]}"
  }
}
```

Note the following:

- The Lambda function has updated the `sessionAttributes`.
- `dialogAction.type` is set to `Close`, which directs Amazon Lex to not expect a user response.
- `dialogAction.fulfillmentState` is set to `Fulfilled`, indicating that the intent is successfully fulfilled.

The client displays the message: **Okay, I have booked your appointment. We will see you at 4:00 p.m. on 2017-02-15.**

Book Trip

This example illustrates creating a bot that is configured to support multiple intents. The example also illustrates how you can use session attributes for cross-intent information sharing. After creating the bot, you use a test client in the Amazon Lex console to test the bot (BookTrip). The client uses the [PostText](#) runtime API operation to send requests to Amazon Lex for each user input.

The BookTrip bot in this example is configured with two intents (BookHotel and BookCar). For example, suppose a user first books a hotel. During the interaction, the user provides information such as check-in dates, location, and number of nights. After the intent is fulfilled, the client can persist this information using session attributes. For more information about session attributes, see [PostText](#).

Now suppose that the user continues to book a car. Using information that the user provided in the previous BookHotel intent (that is, destination city, and check-in and check-out dates), the code hook (Lambda function) you configured to initialize and validate the BookCar intent, initializes slot data for the BookCar intent (that is, destination, pick-up city, pick-up date, and return date). This illustrates how cross-intent information sharing enables you to build bots that can engage in dynamic conversation with the user.

In this example, we use the following session attributes. Only the client and the Lambda function can set and update session attributes. Amazon Lex only passes these between the client and the Lambda function. Amazon Lex doesn't maintain or modify any session attributes.

- `currentReservation` – Contains slot data for an in-progress reservation and other relevant information. For example, the following is a sample request from the client to Amazon Lex. It shows the `currentReservation` session attribute in the request body.

```
POST /bot/BookTrip/alias/$LATEST/user/wch89kjqcpkds8seny7dly5x3otq68j3/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText":"Chicago",
  "sessionAttributes":{"
    "currentReservation":{"ReservationType\":"Hotel",
                        "Location\":"Moscow",
                        "RoomType\:null,
```

```
}
  \"CheckInDate\":null,
  \"Nights\":null}
}
```

- **lastConfirmedReservation** – Contains similar information for a previous intent, if any. For example, if the user booked a hotel and then is in process of booking a car, this session attribute stores slot data for the previous BookHotel intent.
- **confirmationContext** – The Lambda function sets this to AutoPopulate when it prepopulates some of the slot data based on slot data from the previous reservation (if there is one). This enables cross-intent information sharing. For example, if the user previously booked a hotel and now wants to book a car, Amazon Lex can prompt the user to confirm (or deny) that the car is being booked for the same city and dates as their hotel reservation

In this exercise you use blueprints to create an Amazon Lex bot and a Lambda function. For more information about blueprints, see [Amazon Lex and AWS Lambda Blueprints](#).

Next Step

[Step 1: Review the Blueprints Used in this Exercise](#)

Step 1: Review the Blueprints Used in this Exercise

Topics

- [Overview of the Bot Blueprint \(BookTrip\)](#)
- [Overview of the Lambda Function Blueprint \(lex-book-trip-python\)](#)

Overview of the Bot Blueprint (BookTrip)

The blueprint (**BookTrip**) you use to create a bot provides the following preconfiguration:

- **Slot types** – Two custom slot types:

- RoomTypes with enumeration values: king, queen, and deluxe, for use in the BookHotel intent.
- CarTypes with enumeration values: economy, standard, midsize, full size, luxury, and minivan, for use in the BookCar intent.
- **Intent 1 (BookHotel)** – It is preconfigured as follows:
 - **Preconfigured slots**
 - RoomType, of the RoomTypes custom slot type
 - Location, of the AMAZON.US_CITY built-in slot type
 - CheckInDate, of the AMAZON.DATE built-in slot type
 - Nights, of the AMAZON.NUMBER built-in slot type
 - **Preconfigured utterances**
 - "Book a hotel"
 - "I want to make hotel reservations"
 - "Book a {Nights} stay in {Location}"

If the user utters any of these, Amazon Lex determines that BookHotel is the intent and then prompts the user for slot data.

- **Preconfigured prompts**
 - Prompt for the Location slot – "What city will you be staying in?"
 - Prompt for the CheckInDate slot – "What day do you want to check in?"
 - Prompt for the Nights slot – "How many nights will you be staying?"
 - Prompt for the RoomType slot – "What type of room would you like, queen, king, or deluxe?"
 - Confirmation statement – "Okay, I have you down for a {Nights} night stay in {Location} starting {CheckInDate}. Shall I book the reservation?"
 - Denial – "Okay, I have cancelled your reservation in progress."
- **Intent 2 (BookCar)** – It is preconfigured as follows:
 - **Preconfigured slots**
 - PickupCity, of the AMAZON.US_CITY built-in type
 - PickupDate, of the AMAZON.DATE built-in type

- `ReturnDate`, of the `AMAZON.DATE` built-in type
- `DriverAge`, of the `AMAZON.NUMBER` built-in type
- `CarType`, of the `CarTypes` custom type
- **Preconfigured utterances**
 - "Book a car"
 - "Reserve a car"
 - "Make a car reservation"

If the user utters any of these, Amazon Lex determines `BookCar` is the intent and then prompts the user for slot data.

- **Preconfigured prompts**
 - Prompt for the `PickUpCity` slot – "In what city do you need to rent a car?"
 - Prompt for the `PickUpDate` slot – "What day do you want to start your rental?"
 - Prompt for the `ReturnDate` slot – "What day do you want to return this car?"
 - Prompt for the `DriverAge` slot – "How old is the driver for this rental?"
 - Prompt for the `CarType` slot – "What type of car would you like to rent? Our most popular options are economy, midsize, and luxury"
 - Confirmation statement – "Okay, I have you down for a {`CarType`} rental in {`PickUpCity`} from {`PickUpDate`} to {`ReturnDate`}. Should I book the reservation?"
 - Denial – "Okay, I have cancelled your reservation in progress."

Overview of the Lambda Function Blueprint (`lex-book-trip-python`)

In addition to the bot blueprint, AWS Lambda provides a blueprint (`lex-book-trip-python`) that you can use as a code hook with the bot blueprint. For a list of bot blueprints and corresponding Lambda function blueprints, see [Amazon Lex and AWS Lambda Blueprints](#).

When you create a bot using the `BookTrip` blueprint, you update configuration of both the intents (`BookCar` and `BookHotel`) by adding this Lambda function as a code hook for both initialization/validation of user data input and fulfillment of the intents.

This Lambda function code provided showcases dynamic conversation using previously known information (persisted in session attributes) about a user to initialize slot values for an intent. For more information, see [Managing Conversation Context](#).

Next Step

[Step 2: Create an Amazon Lex Bot](#)

Step 2: Create an Amazon Lex Bot

In this section, you create an Amazon Lex bot (BookTrip).

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. On the **Bots** page, choose **Create**.
3. On the **Create your Lex bot** page,
 - Choose **BookTrip** blueprint.
 - Leave the default bot name (BookTrip).
4. Choose **Create**. The console sends a series of requests to Amazon Lex to create the bot. Note the following:
5. The console shows the BookTrip bot. On the **Editor** tab, review the details of the preconfigured intents (BookCar and BookHotel).
6. Test the bot in the test window. Use the following to engage in a test conversation with your bot:

The screenshot shows a chat interface titled "Test Bot". The conversation proceeds as follows:

- User: "Book a hotel"
- Bot: "What city will you be staying in?"
- User: "Chicago"
- Bot: "What day do you want to check in?"
- User: "December 18th"
- Bot: "How many nights will you be staying?"
- User: "4"
- Bot: "What type of room would you like, queen, king or deluxe?"
- User: "Queen"
- Bot: "Okay, I have you down for a 4 night stay in Chicago starting 2016-12-18. Shall I book the reservation?"
- User: "Yes"
- Bot: "CheckInDate:2016-12-18 Location:Chicago Nights:4 RoomType:queen"

Below the final bot message is a "Clear" button. At the bottom of the interface is a text input field with the placeholder "Type to your bot..." and a send button.

From the initial user input ("Book a hotel"), Amazon Lex infers the intent (BookHotel). The bot then uses the prompts preconfigured in this intent to elicit slot data from the user. After user provide all of the slot data, Amazon Lex returns a response back to the client with a message that includes all the user input as a message. The client displays the message in the response as shown.

```
CheckInDate:2016-12-18 Location:Chicago Nights:5 RoomType:queen
```

Now you continue the conversation and try to book a car in the following conversation.

The screenshot shows the 'Test Bot' interface with a blue header bar containing a speech bubble icon, the text 'Test Bot', and a dropdown arrow. The conversation history is as follows:

- User: Also book a car
- Bot: In what city do you need to rent a car?
- User: Chicago
- Bot: What day do you want to start your rental?
- User: December 18th
- Bot: What day do you want to return the car?
- User: December 22nd
- Bot: How old is the driver for this rental?
- User: 35
- Bot: What type of car would you like to rent? Our most popular options are economy, midsize, and luxury
- User: economy
- Bot: The price of this economy rental in Chicago from 2016-12-18 to 2016-12-22 is 556 dollars. Shall I book the reservation?
- User: Yes
- Bot: CarType:economy DriverAge:35 PickUpCity:Chicago
PickUpDate:2016-12-18 ReturnDate:2016-12-22

Below the last bot message is a 'Clear' button. At the bottom of the interface is a text input field with the placeholder text 'Type to your bot...'.

Note that,

- There is no user data validation at this time. For example, you can provide any city to book a hotel.

- You are providing some of the same information again (destination, pick-up city, pick-up date, and return date) to book a car. In a dynamic conversation, your bot should initialize some of this information based on prior input user provided for booking hotel.

In this next section, you create a Lambda function to do some of the user data validation, and initialization using cross-intent information sharing via session attributes. Then you update the intent configuration by adding the Lambda function as code hook to perform initialization/validation of user input and fulfill intent.

Next Step

[Step 3: Create a Lambda function](#)

Step 3: Create a Lambda function

In this section you create a Lambda function using a blueprint (**lex-book-trip-python**) provided in the AWS Lambda console. You also test the Lambda function by invoking it using sample event data provided by the console.

This Lambda function is written in Python.

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. Choose **Create function**.
3. Choose **Use a blueprint**. Type **lex** to find the blueprint, choose the **lex-book-trip-python** blueprint.
4. Choose **Configure** the Lambda function as follows.
 - Type a Lambda function name (**BookTripCodeHook**).
 - For the role, choose **Create a new role from template(s)** and then type a role name.
 - Leave the other default values.
5. Choose **Create function**.
6. If you are using a locale other than English (US) (en-US), update the intent names as described in [Updating a Blueprint for a Specific Locale](#).
7. Test the Lambda function. You invoke the Lambda function twice, using sample data for both booking a car and booking a hotel.

- a. Choose **Configure test event** from the **Select a test event** drop down.
- b. Choose **Amazon Lex Book Hotel** from the **Sample event template** list.

This sample event matches the Amazon Lex request/response model. For more information, see [Using Lambda Functions](#).

- c. Choose **Save and test**.
- d. Verify that the Lambda function ran successfully. The response in this case matches the Amazon Lex response model.
- e. Repeat the step. This time you choose the **Amazon Lex Book Car** from the **Sample event template** list. The Lambda function processes the car reservation.

Next Step


[Step 4: Add the Lambda Function as a Code Hook](#)

Step 4: Add the Lambda Function as a Code Hook

In this section, you update the configurations of both the BookCar and BookHotel intents by adding the Lambda function as a code hook for initialization/validation and fulfillment activities. Make sure you choose the \$LATEST version of the intents because you can only update the \$LATEST version of your Amazon Lex resources.


1. In the Amazon Lex console, choose the **BookTrip** bot.
2. On the **Editor** tab, choose the **BookHotel** intent. Update the intent configuration as follows:
 - a. Make sure the intent version (next to the intent name) is \$LATEST.
 - b. Add the Lambda function as an initialization and validation code hook as follows:
 - In **Options**, choose **Initialization and validation code hook**.
 - Choose your Lambda function from the list.
 - c. Add the Lambda function as a fulfillment code hook as follows:
 - In **Fulfillment**, choose **AWS Lambda function**.
 - Choose your Lambda function from the list.

- Choose **Goodbye message** and type a message.
- d. Choose **Save**.
3. On the **Editor** tab, choose the BookCar intent. Follow the preceding step to add your Lambda function as validation and fulfillment code hook.
 4. Choose **Build**. The console sends a series of requests to Amazon Lex to save the configurations.
 5. Test the bot. Now that you have a Lambda function performing the initialization, user data validation and fulfillment, you can see the difference in the user interaction in the following conversation:



For more information about the data flow from the client (console) to Amazon Lex, and from Amazon Lex to the Lambda function, see [Data Flow: Book Hotel Intent](#).

6. Continue the conversation and book a car as shown in the following image:



When you choose to book a car, the client (console) sends a request to Amazon Lex that includes the session attributes (from the previous conversation, BookHotel). Amazon Lex passes this information to the Lambda function, which then initializes (that is, it prepopulates) some of the BookCar slot data (that is, PickUpDate, ReturnDate, and PickUpCity).

Note

This illustrates how session attributes can be used to maintain context across intents. The console client provides the **Clear** link in the test window that a user can use to clear any prior session attributes.

For more information about the data flow from the client (console) to Amazon Lex, and from Amazon Lex to the Lambda function, see [Data Flow: Book Car Intent](#).

Details of the Information Flow

In this exercise, you engaged in a conversation with the Amazon Lex BookTrip bot using the test window client provided in the Amazon Lex console. This section explains the following:

- The data flow between the client and Amazon Lex.

The section assumes that the client sends requests to Amazon Lex using the PostText runtime API and shows request and response details accordingly. For more information about the PostText runtime API, see [PostText](#).

Note

For an example of the information flow between the client and Amazon Lex in which the client uses the PostContent API, see [Step 2a \(Optional\): Review the Details of the Spoken Information Flow \(Console\)](#).

- The data flow between Amazon Lex and the Lambda function. For more information, see [Lambda Function Input Event and Response Format](#).

Topics

- [Data Flow: Book Hotel Intent](#)

- [Data Flow: Book Car Intent](#)

Data Flow: Book Hotel Intent

This section explains what happens after each user input.

1. User: "book a hotel"

- a. The client (console) sends the following [PostText](#) request to Amazon Lex:

```
POST /bot/BookTrip/alias/$LATEST/user/wch89kjqcpkds8seny7dly5x3otq68j3/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText":"book a hotel",
  "sessionAttributes":{}
}
```

Both the request URI and the body provides information to Amazon Lex:

- Request URI – Provides bot name (BookTrip), bot alias (\$LATEST) and the user name. The trailing text indicates that it is a PostText API request (and not PostContent).
 - Request body – Includes the user input (inputText) and empty sessionAttributes. Initially, this is an empty object and the Lambda function first sets the session attributes.
- b. From the inputText, Amazon Lex detects the intent (BookHotel). This intent is configured with a Lambda function as a code hook for user data initialization/validation. Therefore, Amazon Lex invokes that Lambda function by passing the following information as the event parameter (see [Input Event Format](#)):

```
{
  "messageVersion":"1.0",
  "invocationSource":"DialogCodeHook",
  "userId":"wch89kjqcpkds8seny7dly5x3otq68j3",
  "sessionAttributes":{
  },
  "bot":{
```

```

    "name": "BookTrip",
    "alias": null,
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "BookHotel",
    "slots": {
      "RoomType": null,
      "CheckInDate": null,
      "Nights": null,
      "Location": null
    },
    "confirmationStatus": "None"
  }
}

```

In addition to the information sent by the client, Amazon Lex also includes the following additional data:

- **messageVersion** – Currently Amazon Lex supports only the 1.0 version.
 - **invocationSource** – Indicates the purpose of Lambda function invocation. In this case, it is to perform user data initialization and validation (at this time Amazon Lex knows that the user has not provided all the slot data to fulfill the intent).
 - **currentIntent** – All of the slot values are set to null.
- c. At this time, all the slot values are null. There is nothing for the Lambda function to validate. The Lambda function returns the following response to Amazon Lex. For information about response format, see [Response Format](#).

```

{
  "sessionAttributes": {
    "currentReservation": "{\"ReservationType\": \"Hotel\", \"Location\": null, \"RoomType\": null, \"CheckInDate\": null, \"Nights\": null}"
  },
  "dialogAction": {
    "type": "Delegate",
    "slots": {
      "RoomType": null,
      "CheckInDate": null,
      "Nights": null,
      "Location": null
    }
  }
}

```



```

    }
  }
}

```

Note


- `currentReservation` – The Lambda function includes this session attribute. Its value is a copy of the current slot information and the reservation type.

Only the Lambda function and the client can update these session attributes. Amazon Lex simply passes these values.

- `dialogAction.type` – By setting this value to `Delegate`, the Lambda function delegates the responsibility for the next course of action to Amazon Lex.

If the Lambda function detected anything in the user data validation, it instructs Amazon Lex what to do next.

- As per the `dialogAction.type`, Amazon Lex decides the next course of action—elicit data from the user for the `Location` slot. It selects one of the prompt messages ("What city will you be staying in?") for this slot, according to the intent configuration, and then sends the following response to the user:



The session attributes are passed to the client.

The client reads the response and then displays the message: "What city will you be staying in?"

2. User: "Moscow"

- a. The client sends the following PostText request to Amazon Lex (line breaks added for readability):

```
POST /bot/BookTrip/alias/$LATEST/user/wch89kjqcpkds8seny7dly5x3otq68j3/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText":"Moscow",
  "sessionAttributes":{"
    "currentReservation":{"\ "ReservationType\":"\ "Hotel\","
                        \ "Location\":null,
                        \ "RoomType\":null,
                        \ "CheckInDate\":null,
                        \ "Nights\":null}"
  }
}
```

In addition to the `inputText`, the client includes the same `currentReservation` session attributes it received.

- b. Amazon Lex first interprets the `inputText` in the context of the current intent (the service remembers that it had asked the specific user for information about `Location` slot). It updates the slot value for the current intent and invokes the Lambda function using the following event:

```
{
  "messageVersion": "1.0",
  "invocationSource": "DialogCodeHook",
  "userId": "wch89kjqcpkds8seny7dly5x3otq68j3",
  "sessionAttributes": {
    "currentReservation": "{\ "ReservationType\":"\ "Hotel\","Location
\ ":null,\ "RoomType\":null,\ "CheckInDate\":null,\ "Nights\":null}"
  },
  "bot": {
    "name": "BookTrip",
    "alias": null,
```

```
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "BookHotel",
    "slots": {
      "RoomType": null,
      "CheckInDate": null,
      "Nights": null,
      "Location": "Moscow"
    },
    "confirmationStatus": "None"
  }
}
```

Note

- `invocationSource` continues to be `DialogCodeHook`. In this step, we are just validating user data.
- Amazon Lex is just passing the session attribute to the Lambda function.
- For `currentIntent.slots`, Amazon Lex has updated the `Location` slot to `Moscow`.

- c. The Lambda function performs the user data validation and determines that `Moscow` is an invalid location.

Note

The Lambda function in this exercise has a simple list of valid cities and `Moscow` is not on the list. In a production application, you might use a back-end database to get this information.

It resets the slot value back to null and directs Amazon Lex to prompt the user again for another value by sending the following response:


```
{
  "sessionAttributes": {
```

```
    "currentReservation": "{\\"ReservationType\\":\\"Hotel\\",\\"Location\\":\\"Moscow\\",\\"RoomType\\":null,\\"CheckInDate\\":null,\\"Nights\\":null}"
  },
  "dialogAction": {
    "type": "ElicitSlot",
    "intentName": "BookHotel",
    "slots": {
      "RoomType": null,
      "CheckInDate": null,
      "Nights": null,
      "Location": null
    },
    "slotToElicit": "Location",
    "message": {
      "contentType": "PlainText",
      "content": "We currently do not support Moscow as a valid destination. Can you try a different city?"
    }
  }
}
```

 Note

- `currentIntent.slots.Location` is reset to null.
- `dialogAction.type` is set to `ElicitSlot`, which directs Amazon Lex to prompt the user again by providing the following:
 - `dialogAction.slotToElicit` – slot for which to elicit data from the user.
 - `dialogAction.message` – a message to convey to the user.

- d. Amazon Lex notices the `dialogAction.type` and passes the information to the client in the following response:



The client simply displays the message: "We currently do not support Moscow as a valid destination. Can you try a different city?"

3. User: "Chicago"

- a. The client sends the following PostText request to Amazon Lex:

```

POST /bot/BookTrip/alias/$LATEST/user/wch89kjqcpkds8seny7dly5x3otq68j3/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText":"Chicago",
  "sessionAttributes":{
    "currentReservation":{"ReservationType\":"Hotel\",
                        \"Location\":"Moscow\",
                        \"RoomType\":null,
                        \"CheckInDate\":null,
                        \"Nights\":null}}
  }
}

```

- b. Amazon Lex knows the context, that it was eliciting data for the Location slot. In this context, it knows the inputText value is for the Location slot. It then invokes the Lambda function by sending the following event:

```
{
  "messageVersion": "1.0",
  "invocationSource": "DialogCodeHook",
  "userId": "wch89kjqcpkds8seny7dly5x3otq68j3",
  "sessionAttributes": {
    "currentReservation": "{\"ReservationType\": \"Hotel\", \"Location\": \"Moscow\", \"RoomType\": null, \"CheckInDate\": null, \"Nights\": null}"
  },
  "bot": {
    "name": "BookTrip",
    "alias": null,
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "BookHotel",
    "slots": {
      "RoomType": null,
      "CheckInDate": null,
      "Nights": null,
      "Location": "Chicago"
    },
    "confirmationStatus": "None"
  }
}
```

Amazon Lex updated the `currentIntent.slots` by setting the `Location` slot to `Chicago`.

- c. According to the `invocationSource` value of `DialogCodeHook`, the Lambda function performs user data validation. It recognizes `Chicago` as a valid slot value, updates the session attribute accordingly, and then returns the following response to Amazon Lex.


```
{
  "sessionAttributes": {
    "currentReservation": "{\"ReservationType\": \"Hotel\", \"Location\": \"Chicago\", \"RoomType\": null, \"CheckInDate\": null, \"Nights\": null}"
  },
  "dialogAction": {
    "type": "Delegate",
    "slots": {
      "RoomType": null,
```

```
        "CheckInDate": null,  
        "Nights": null,  
        "Location": "Chicago"  
    }  
}
```

Note

- `currentReservation` – The Lambda function updates this session attribute by setting the `Location` to `Chicago`.
- `dialogAction.type` – Is set to `Delegate`. User data was valid, and the Lambda function directs Amazon Lex to choose the next course of action.

- d. According to `dialogAction.type`, Amazon Lex chooses the next course of action. Amazon Lex knows that it needs more slot data and picks the next unfilled slot (`CheckInDate`) with the highest priority according to the intent configuration. It selects one of the prompt messages ("What day do you want to check in?") for this slot according to the intent configuration and then sends the following response back to the client:




The client displays the message: "What day do you want to check in?"

4. The user interaction continues—the user provides data, the Lambda function validates data, and then delegates the next course of action to Amazon Lex. Eventually the user provides all of the slot data, the Lambda function validates all of the user input, and then Amazon Lex recognizes it has all the slot data.

Note

In this exercise, after the user provides all of the slot data, the Lambda function computes the price of the hotel reservation and returns it as another session attribute (`currentReservationPrice`).

At this point, the intent is ready to be fulfilled, but the `BookHotel` intent is configured with a confirmation prompt requiring user confirmation before Amazon Lex can fulfill the intent. Therefore, Amazon Lex sends the following message to the client requesting confirmation before booking the hotel:



The client display the message: "Okay, I have you down for a 5 night in Chicago starting 2016-12-18. Shall I book the reservation?"

5. User: "yes"
 - a. The client sends the following `PostText` request to Amazon Lex:

```
POST /bot/BookTrip/alias/$LATEST/user/wch89kjqcpkds8seny7dly5x3otq68j3/text
"Content-Type": "application/json"
```



```
"Content-Encoding": "amz-1.0"

{
  "inputText": "Yes",
  "sessionAttributes": {
    "currentReservation": "{\\"ReservationType\\":\\"Hotel\\",
                          \\"Location\\":\\"Chicago\\",
                          \\"RoomType\\":\\"queen\\",
                          \\"CheckInDate\\":\\"2016-12-18\\",
                          \\"Nights\\":\\"5\\"}",
    "currentReservationPrice": "1195"
  }
}
```

- b. Amazon Lex interprets the `inputText` in the context of confirming the current intent. Amazon Lex understands that the user wants to proceed with the reservation. This time Amazon Lex invokes the Lambda function to fulfill the intent by sending the following event. By setting the `invocationSource` to `FulfillmentCodeHook` in the event, it sends to the Lambda function. Amazon Lex also sets the `confirmationStatus` to `Confirmed`.

```
{
  "messageVersion": "1.0",
  "invocationSource": "FulfillmentCodeHook",
  "userId": "wch89kjqcpkds8seny7dly5x3otq68j3",
  "sessionAttributes": {
    "currentReservation": "{\\"ReservationType\\":\\"Hotel\\",\\"Location\\":
    \\"Chicago\\",\\"RoomType\\":\\"queen\\",\\"CheckInDate\\":\\"2016-12-18\\",\\"Nights\\":
    \\"5\\"}",
    "currentReservationPrice": "956"
  },
  "bot": {
    "name": "BookTrip",
    "alias": null,
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "BookHotel",
    "slots": {
      "RoomType": "queen",
      "CheckInDate": "2016-12-18",
      "Nights": "5",

```

```

        "Location": "Chicago"
      },
      "confirmationStatus": "Confirmed"
    }
  }
}

```

Note

- `invocationSource` – This time, Amazon Lex set this value to `FulfillmentCodeHook`, directing the Lambda function to fulfill the intent.
- `confirmationStatus` – Is set to `Confirmed`.

- c. This time, the Lambda function fulfills the `BookHotel` intent, Amazon Lex completes the reservation, and then it returns the following response:

```

{
  "sessionAttributes": {
    "lastConfirmedReservation": "{\"ReservationType\":\"Hotel\",\"Location\":\"Chicago\",\"RoomType\":\"queen\",\"CheckInDate\":\"2016-12-18\",\"Nights\":\"5\"}"
  },
  "dialogAction": {
    "type": "Close",
    "fulfillmentState": "Fulfilled",
    "message": {
      "contentType": "PlainText",
      "content": "Thanks, I have placed your reservation. Please let me know if you would like to book a car rental, or another hotel."
    }
  }
}


```

Note

- `lastConfirmedReservation` – Is a new session attribute that the Lambda function added (instead of the `currentReservation`, `currentReservationPrice`).
- `dialogAction.type` – The Lambda function sets this value to `Close`, indicating that Amazon Lex to not expect a user response.

- `dialogAction.fulfillmentState` – Is set to `Fulfilled` and includes an appropriate message to convey to the user.

- d. Amazon Lex reviews the `fulfillmentState` and sends the following response to the client:



Note

- `dialogState` – Amazon Lex sets this value to `Fulfilled`.
- `message` – Is the same message that the Lambda function provided.

The client displays the message.

Data Flow: Book Car Intent

The BookTrip bot in this exercise supports two intents (`BookHotel` and `BookCar`). After booking a hotel, the user can continue the conversation to book a car. As long as the session hasn't timed out, in each subsequent request the client continues to send the session attributes (in this example, the `lastConfirmedReservation`). The Lambda function can use this information to initialize slot data for the `BookCar` intent. This shows how you can use session attributes in cross-intent data sharing.

Specifically, when the user chooses the BookCar intent, the Lambda function uses relevant information in the session attribute to prepopulate slots (PickUpDate, ReturnDate, and PickUpCity) for the BookCar intent.

 Note

The Amazon Lex console provides the **Clear** link that you can use to clear any prior session attributes.

Follow the steps in this procedure to continue the conversation.

1. User: "also book a car"

- a. The client sends the following PostText request to Amazon Lex.

```
POST /bot/BookTrip/alias/$LATEST/user/wch89kjqcpkds8seny7dly5x3otq68j3/text
"Content-Type":"application/json"
"Content-Encoding":"amz-1.0"

{
  "inputText":"also book a car",
  "sessionAttributes":{
    "lastConfirmedReservation":""{"ReservationType\":"Hotel\","
                                \Location\":"Chicago\","
                                \RoomType\":"queen\","
                                \CheckInDate\":"2016-12-18\","
                                \Nights\":"5\"}
  }
}
```

The client includes the lastConfirmedReservation session attribute.

- b. Amazon Lex detects the intent (BookCar) from the inputText. This intent is also configured to invoke the Lambda function to perform the initialization and validation of the user data. Amazon Lex invokes the Lambda function with the following event:

```
{
  "messageVersion": "1.0",
  "invocationSource": "DialogCodeHook",
  "userId": "wch89kjqcpkds8seny7dly5x3otq68j3",
  "sessionAttributes": {
```

```

    "lastConfirmedReservation": "{\"ReservationType\":\"Hotel\",\"Location\":"
    "\":\"Chicago\",\"RoomType\":\"queen\",\"CheckInDate\":\"2016-12-18\",
    \"Nights\":\"5\"}"
  },
  "bot": {
    "name": "BookTrip",
    "alias": null,
    "version": "$LATEST"
  },
  "outputDialogMode": "Text",
  "currentIntent": {
    "name": "BookCar",
    "slots": {
      "PickUpDate": null,
      "ReturnDate": null,
      "DriverAge": null,
      "CarType": null,
      "PickUpCity": null
    },
    "confirmationStatus": "None"
  }
}

```

Note

- `messageVersion` – Currently Amazon Lex supports the 1.0 version only.
- `invocationSource` – Indicates the purpose of invocation is to perform initialization and user data validation.
- `currentIntent` – It includes the intent name and the slots. At this time, all slot values are null.

- c. The Lambda function notices all null slot values with nothing to validate. However, it uses session attributes to initialize some of the slot values (`PickUpDate`, `ReturnDate`, and `PickUpCity`), and then returns the following response:

```

{
  "sessionAttributes": {

```

```

    "lastConfirmedReservation": "{\"ReservationType\":\"Hotel\",\"Location\":"
    "\":\"Chicago\",\"RoomType\":\"queen\",\"CheckInDate\":\"2016-12-18\",\"Nights\":"
    "\":\"5\"}",
    "currentReservation": "{\"ReservationType\":\"Car\",\"PickUpCity\":"
    "\":null,\"PickUpDate\":null,\"ReturnDate\":null,\"CarType\":null}",
    "confirmationContext": "AutoPopulate"
  },
  "dialogAction": {
    "type": "ConfirmIntent",
    "intentName": "BookCar",
    "slots": {
      "PickUpCity": "Chicago",
      "PickUpDate": "2016-12-18",
      "ReturnDate": "2016-12-22",
      "CarType": null,
      "DriverAge": null
    },
    "message": {
      "contentType": "PlainText",
      "content": "Is this car rental for your 5 night stay in Chicago on
2016-12-18?"
    }
  }
}

```

Note


- In addition to the `lastConfirmedReservation`, the Lambda function includes more session attributes (`currentReservation` and `confirmationContext`).
- `dialogAction.type` is set to `ConfirmIntent`, which informs Amazon Lex that a yes, no reply is expected from the user (the `confirmationContext` set to `AutoPopulate`, the Lambda function knows that the yes/no user reply is to obtain user confirmation of the initialization the Lambda function performed (auto populated slot data).

The Lambda function also includes in the response an informative message in the `dialogAction.message` for Amazon Lex to return to the client.

Note

The term `ConfirmIntent` (value of the `dialogAction.type`) is not related to any bot intent. In the example, Lambda function uses this term to direct Amazon Lex to get a yes/no reply from the user.

- d. According to the `dialogAction.type`, Amazon Lex returns the following response to the client:



The client displays the message: "Is this car rental for your 5 night stay in Chicago on 2016-12-18?"

2. User: "yes"

- a. The client sends the following `PostText` request to Amazon Lex.

```

POST /bot/BookTrip/alias/$LATEST/user/wch89kjqcpkds8seny7dly5x3otq68j3/text
"Content-Type": "application/json"
"Content-Encoding": "amz-1.0"

{
  "inputText": "yes",
  "sessionAttributes": {
    "confirmationContext": "AutoPopulate",
    "currentReservation": {
      "ReservationType": "Car",

```

```

        \"PickUpCity\":null,
        \"PickUpDate\":null,
        \"ReturnDate\":null,
        \"CarType\":null}],
    \"lastConfirmedReservation\":{\"ReservationType\":\"Hotel\",
        \"Location\":\"Chicago\",
        \"RoomType\":\"queen\",
        \"CheckInDate\":\"2016-12-18\",
        \"Nights\":\"5\"}
  }
}

```

- b. Amazon Lex reads the `inputText` and it knows the context (asked the user to confirm the auto population). Amazon Lex invokes the Lambda function by sending the following event:

```

{
  \"messageVersion\": \"1.0\",
  \"invocationSource\": \"DialogCodeHook\",
  \"userId\": \"wch89kjqcpkds8seny7dly5x3otq68j3\",
  \"sessionAttributes\": {
    \"confirmationContext\": \"AutoPopulate\",
    \"currentReservation\": \"{\\\"ReservationType\\\":\\\"Car\\\",\\\"PickUpCity\\\":null,\\\"PickUpDate\\\":null,\\\"ReturnDate\\\":null,\\\"CarType\\\":null}\",
    \"lastConfirmedReservation\": \"{\\\"ReservationType\\\":\\\"Hotel\\\",\\\"Location\\\":\\\"Chicago\\\",\\\"RoomType\\\":\\\"queen\\\",\\\"CheckInDate\\\":\\\"2016-12-18\\\",\\\"Nights\\\":\\\"5\\\"}\"
  },
  \"bot\": {
    \"name\": \"BookTrip\",
    \"alias\": null,
    \"version\": \"$LATEST\"
  },
  \"outputDialogMode\": \"Text\",
  \"currentIntent\": {
    \"name\": \"BookCar\",
    \"slots\": {
      \"PickUpDate\": \"2016-12-18\",
      \"ReturnDate\": \"2016-12-22\",
      \"DriverAge\": null,
      \"CarType\": null,
      \"PickUpCity\": \"Chicago\"
    }
  },
}

```



```

        "confirmationStatus": "Confirmed"
    }
}

```

Because the user replied Yes, Amazon Lex sets the `confirmationStatus` to `Confirmed`.

- c. From the `confirmationStatus`, the Lambda function knows that the prepopulated values are correct. The Lambda function does the following:
 - Updates the `currentReservation` session attribute to slot value it had prepopulated.
 - Sets the `dialogAction.type` to `ElicitSlot`
 - Sets the `slotToElicit` value to `DriverAge`.

The following response is sent:


```

{
  "sessionAttributes": {
    "currentReservation": "{\"ReservationType\":\"Car\",\"PickUpCity\": \"Chicago\", \"PickUpDate\": \"2016-12-18\", \"ReturnDate\": \"2016-12-22\", \"CarType\": null}\",
    "lastConfirmedReservation": "{\"ReservationType\":\"Hotel\", \"Location\": \"Chicago\", \"RoomType\": \"queen\", \"CheckInDate\": \"2016-12-18\", \"Nights\": \"5\"}"
  },
  "dialogAction": {
    "type": "ElicitSlot",
    "intentName": "BookCar",
    "slots": {
      "PickUpDate": "2016-12-18",
      "ReturnDate": "2016-12-22",
      "DriverAge": null,
      "CarType": null,
      "PickUpCity": "Chicago"
    },
    "slotToElicit": "DriverAge",
    "message": {
      "contentType": "PlainText",
      "content": "How old is the driver of this car rental?"
    }
  }
}

```

```
}
```

d. Amazon Lex returns following response:




The client displays the message "How old is the driver of this car rental?" and the conversation continues.

Using a Response Card

In this exercise, you extend Getting Started Exercise 1 by adding a response card. You create a bot that supports the OrderFlowers intent, and then update the intent by adding a response card for the FlowerType slot. In addition to the following prompt for the FlowerType slot, the user can choose the type of flowers from the response card:

```
What type of flowers would you like to order?
```

The following is the response card:



The bot user can either type the text or choose from the list of flower types. This response card is configured with an image, which appears in the client as shown. For more information about response cards, see [Response Cards](#).

To create and test a bot with a response card:

1. Follow Getting Started Exercise 1 to create and test an OrderFlowers bot. You must complete steps 1, 2, and 3. You don't need to add a Lambda function to test the response card. For instructions, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#).
2. Update the bot by adding the response card, and then publish a version. When you publish a version, specify an alias (BETA) to point to it.
 - a. In the Amazon Lex console, choose your bot.
 - b. Choose the OrderFlowers intent.
 - c. Choose the settings gear icon next to the "What type of flowers" **Prompt** to configure a response card for the FlowerType, as shown in the following image.

The screenshot shows the Amazon Lex console interface for configuring an intent named "OrderFlowers". The interface includes tabs for Editor, Settings, Channels, and Monitoring. The Editor tab is active, showing sample utterances and a slots table. The slots table has columns for Priority, Required, Name, Slot type, and Prompt. Three slots are configured: FlowerType, PickupDate, and PickupTime. The first slot, FlowerType, is highlighted with a red box around its settings icon.

Priority	Required	Name	Slot type	Prompt
		e.g. Location	e.g. A...	e.g. What city?
1.	✓	FlowerType	Flower...	Latest
2.	✓	PickupDate	AMAZ...	Built-in
3.	✓	PickupTime	AMAZ...	Built-in

- d. Give the card a title and configure three buttons as shown in the following screen shot. You can optionally add an image to the response card, provided you have an image URL. If you are deploying your bot using Twilio SMS, you must provide an image URL.

Prompt response cards

Card 1 ⓘ

Preview as: Facebook ▼

🗑️

Image URL*

e.g. http://www.example.com/image.png

Title*

What type of flowers?

Subtitle*

Button title*

Tulips

✕

Button value*

tulips ▼

Button title

Lilies

✕

Button value

lilies ▼

Button title

Roses

✕

Button value

roses ▼

+

Add Card

- e. Choose **Save** to save the response card.
- f. Choose **Save intent** to save the intent configuration.
- g. To build the bot, choose **Build**.
- h. To publish a bot version, choose **Publish**. Specify BETA as an alias that points to the bot version. For information about versioning, see [Versioning and Aliases](#).

3. Deploy the bot on a messaging platform:

- Deploy the bot on the Facebook Messenger platform and test the integration. For instructions, see [Integrating an Amazon Lex Bot with Facebook Messenger](#). When you order flowers, the message window shows the response card so you can choose a flower type.
- Deploy the bot on the Slack platform and test the integration. For instructions, see [Integrating an Amazon Lex Bot with Slack](#). When you order flowers, the message window shows the response card so you can choose a flower type.
- Deploy the bot on the Twilio SMS platform. For instructions, see [Integrating an Amazon Lex Bot with Twilio Programmable SMS](#). When you order flowers, the message from Twilio shows the image from the response card. Twilio SMS does not support buttons in the response.

Updating Utterances

In this exercise, you add additional utterances to those you created in Getting Started Exercise 1. You use the **Monitoring** tab in the Amazon Lex console to view utterances that your bot did not recognize. To improve the experience for your users, you add those utterances to the bot.

Utterance statistics are not generated under the following conditions:

- The `childDirected` field was set to true when the bot was created.
- You are using slot obfuscation with one or more slots.
- You opted out of participating in improving Amazon Lex.

Note

Utterance statistics are generated once a day. You can see the utterance that was not recognized, how many times it was heard, and the last date and time that the utterance was heard. It can take up to 24 hours for missed utterances to appear in the console.

You can see utterances for different versions of your bot. To change the version of your bot that you are seeing utterances for, choose a different version from the drop-down next to the bot name.

To view and add missed utterances to a bot:

1. Follow the first step of Getting Started Exercise 1 to create and test an `OrderFlowers` bot. For instructions, see [Exercise 1: Create an Amazon Lex Bot Using a Blueprint \(Console\)](#).
2. Test the bot by typing the following utterances in the **Test Bot** window. Type each utterance several times. The example bot doesn't recognize the following utterances:
 - Order flowers
 - Get me flowers
 - Please order flowers
 - Get me some flowers
3. Wait for Amazon Lex to gather usage data about the missed utterances. Utterance data is generated once per day, generally overnight.
4. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
5. Choose the `OrderFlowers` bot.
6. Choose the **Monitoring** tab, and then choose **Utterances** from the left menu and then choose the **Missed** button. The following pane shows a maximum of 100 missed utterances.

Utterances				
Add utterance to Intent ▾				
Filter: 🔍 Filter by keyword				
Detected Missed				
☐	Utterances ▾	Count	Status	Last said date ▾
☐	I want flowers	5	Missed	April 21, 2017 at 10:28:13 A
☐	Order flowers	4	Missed	April 21, 2017 at 10:28:05 A
☐	Get me some flowers	2	Missed	April 21, 2017 at 10:27:49 A
☐	Get me flowers	2	Missed	April 21, 2017 at 10:27:25 A
☐	Please order flowers	1	Missed	April 21, 2017 at 10:26:55 A
☐	get me some flowers	1	Missed	April 21, 2017 at 10:27:18 A

7. To choose the missed utterances that you want to add to the bot, select the check box next to them. To add the utterance to the `$LATEST` version of the intent, choose the down arrow next to the **Add utterance to intent** dropdown, and then choose the intent.

8. To rebuild your bot, choose **Build** and then **Build** again to re-build your bot.
9. To verify that your bot recognizes the new utterances, use the **Test Bot** pane.

Integrating with a Web site

In this example you integrate a bot with a Web site using text and voice. You use JavaScript and AWS services to build an interactive experience for visitors to your Web site. You can choose from these examples documented on the [AWS AI Blog](#):

- [Deploy a Web UI for Your Chatbot](#)—Demonstrates a full-featured Web UI that provides a Web client for Amazon Lex chatbots. You can use this to learn about Web clients, or as a building block for your own application.
- ["Greetings, visitor!"—Engage Your Web Users with Amazon Lex](#)—Demonstrates using Amazon Lex, the AWS SDK for JavaScript in the Browser, and Amazon Cognito to create a conversational experience on your Web site.
- [Capturing Voice Input in a Browser and Sending it to Amazon Lex](#)—Demonstrates embedding a voice-based chatbot in a Web site using the SDK for JavaScript in the Browser. The application records audio, sends the audio to Amazon Lex, and then plays the response.


Call Center Agent Assistant

In this tutorial, you use Amazon Lex with Amazon Kendra to build an agent assist bot that assists customer support agents and publish it as a web application. Amazon Kendra is an enterprise search service that uses machine learning to search through documents to find answers. For more information about Amazon Kendra, see [the Amazon Kendra Developer Guide](#).

Amazon Lex bots are widely used in call centers as the first point of contact for customers. A bot is often capable of resolving customer questions. When a bot can't answer a question, it transfers the conversation to a customer support employee.

In this tutorial, we create an Amazon Lex bot that agents use to answer customer queries in real time. By reading the answers that the bot provides, the agent is spared from looking up answers manually.

The bot and web application that you create in this tutorial helps agents respond to customers efficiently and accurately by quickly providing the right resources. The following diagram shows how the web application works.



As the diagram shows, the Amazon Kendra index of documents is stored in an Amazon Simple Storage Service (Amazon S3) bucket. If you don't already have an S3 bucket, you can set one up when you create the Amazon Kendra index. In addition to Amazon S3, you will use Amazon Cognito for this tutorial. Amazon Cognito manages permissions for deploying the bot as a web application.

In this tutorial, you create an Amazon Kendra index that provides answers to customer questions, create the bot and add intents that allow it to suggest answers based on the conversation with the customer, set up Amazon Cognito to manage access permissions, and deploy the bot as a web application.

Estimated time: 75 minutes

Estimated cost: \$2.50 per hour for an Amazon Kendra index and \$0.75 for 1000 Amazon Lex requests. Your Amazon Kendra index continues to run after you are finished with this exercise. Be sure to delete it to avoid unnecessary costs.

Note: Make sure that you choose the same AWS Region for all the services used in this tutorial.

Topics

- [Step 1: Create an Amazon Kendra Index](#)
- [Step 2: Create an Amazon Lex Bot](#)
- [Step 3: Add a Custom and Built-in Intent](#)
- [Step 4: Set up Amazon Cognito](#)
- [Step 5: Deploy Your Bot as a Web Application](#)

- [Step 6: Use the Bot](#)

Step 1: Create an Amazon Kendra Index

Begin by creating an Amazon Kendra index of documents that answer customer questions. An index provides a search API for client queries. You create the index from source documents. Amazon Kendra returns answers it finds in indexed documents to the bot, which displays them to the agent.

The quality and accuracy of the responses suggested by Amazon Kendra depend on the documents that you index. Documents should include files that are frequently accessed by the agent and must be stored in an S3 bucket. You can index unstructured and semi-structured data in .html, Microsoft Office (.doc, .ppt), PDF, and text formats.

To create an Amazon Kendra index, see [Getting started with an S3 bucket \(console\)](#) in the *Amazon Kendra Developer Guide*.

To add questions and answers (FAQs) that help answer customer queries, see [Adding questions and answers](#) in the *Amazon Kendra Developer Guide*. For this tutorial, use the [ML_FAQ.csv file on GitHub](#).

Next step

[Step 2: Create an Amazon Lex Bot](#)

Step 2: Create an Amazon Lex Bot

Amazon Lex provides an interface between the call center agent and the Amazon Kendra index. It keeps track of the conversation between the agent and the customer and calls the `AMAZON.KendraSearchIntent` intent based on the questions the customer asks. An *intent* is an action that the user wants to perform.

Amazon Kendra searches the indexed documents and returns an answer to Amazon Lex that it displays in the bot. This answer is visible only to the agent.

To create an agent assistant bot

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. In the navigation pane, choose **Bots**.
3. Choose **Create**.

4. Choose **Custom bot** and configure the bot.
 - a. **Bot name** – Enter a name that indicates the bot's purpose, such as **AgentAssistBot**.
 - b. **Output voice** – Choose **None**.
 - c. **Session timeout** – Enter **5**.
 - d. **COPPA** – Choose **No**.
5. Choose **Create**. After creating the bot, Amazon Lex displays the bot editor tab.

Next step

[Step 3: Add a Custom and Built-in Intent](#)

Step 3: Add a Custom and Built-in Intent

An *intent* represents an action that the call center agent wants the bot to perform. In this case, the agent wants the bot to suggest responses and helpful resources based on the agent's conversation with the customer.

Amazon Lex has two types of intents: custom intents and built-in intents.

`AMAZON.KendraSearchIntent` is a built-in intent. The bot uses the `AMAZON.KendraSearchIntent` intent to query the index and display the responses suggested by Amazon Kendra.

The bot in this example doesn't need a custom intent. However, to build the bot, you must create at least one custom intent with at least one sample utterance. This intent is required only to build your agent assistant bot. It doesn't perform any other function. The utterance for the intent must not answer any of the questions that the customer might ask. This ensures that the `AMAZON.KendraSearchIntent` is called to answer customer queries. For more information, see [AMAZON.KendraSearchIntent](#).

To create the required custom intent

1. On the **Getting started with your bot** page, choose **Create intent**.
2. For **Add intent**, choose **Create intent**.
3. In the **Create intent** dialog box, enter a descriptive name for the intent, such as **RequiredIntent**.
4. For **Sample utterances**, enter a descriptive utterance, such as **Required utterance**.

5. Choose **Save intent**.

To add the **AMAZON.KendraSearchIntent** intent and response message

1. In the navigation pane, choose the plus sign (+) next to **Intents**.
2. Choose **Search existing intents**.
3. In the **Search intents** box, enter **AMAZON.KendraSearchIntent**, then choose it from the list.
4. Give the intent a descriptive name, such as **AgentAssistSearchIntent**, then choose **Add**.
5. In the intent editor, choose **Amazon Kendra query** to open the query options.
6. Choose the index that you want the intent to search,
7. In the **Response** section, add the following three messages to a message group.

```
I found an answer for the customer query: ((x-amz-lex:kendra-search-response-question_answer-question-1)) and the answer is ((x-amz-lex:kendra-search-response-question_answer-answer-1)).  
I found an excerpt from a helpful document: ((x-amz-lex:kendra-search-response-document-1)).  
I think this answer will help the customer: ((x-amz-lex:kendra-search-response-answer-1)).
```

8. Choose **Save intent**.
9. Choose **Build** to build the bot.

Next step

[Step 4: Set up Amazon Cognito](#)

Step 4: Set up Amazon Cognito

To manage permissions and users for the web application, you need to set up Amazon Cognito. Amazon Cognito ensures that the web application is secure and has access control. Amazon Cognito uses identity pools to provide AWS credentials that grant your users access to other AWS services. For this tutorial, it provides access to Amazon Lex.

When creating an identity pool, Amazon Cognito provides you with AWS Identity and Access Management (IAM) roles for authenticated and unauthenticated users. You modify the IAM roles by adding policies that grant access to Amazon Lex.

To set up Amazon Cognito

1. Sign into the AWS Management Console and open the Amazon Cognito console at <https://console.aws.amazon.com/cognito/>.
2. Choose **Manage Identity Pools**.
3. Choose **Create new identity pool**.
4. Configure the identity pool.
 - a. **Identity pool name** – Enter a name that indicates the pool's purpose, such as **BotPool1**.
 - b. In the **Unauthenticated identities** section, choose **Enable access to unauthenticated identities**.
5. Choose **Create Pool**.
6. On the **Identify the IAM roles to use with your new identity pool** page, choose **View Details**.
7. Record the IAM role names. You will modify them later.
8. Choose **Allow**.
9. On the **Getting Started with Amazon Cognito** page, for **Platform**, choose **JavaScript**.
10. In the **Get AWS Credentials** section, find and record the **Identity pool ID**.
11. To allow access to Amazon Lex, modify the authenticated and unauthenticated IAM roles.
 - a. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
 - b. In the navigation pane, under **Access Management**, choose **Roles**.
 - c. In the search box, enter the name of the authenticated IAM role and choose the checkbox next to it.
 - i. Choose **Attach policies**.
 - ii. In the search box, enter **AmazonLexRunBotsOnly** and choose the checkbox next to it.
 - iii. Choose **Attach policy**.
 - d. Enter the name of the unauthenticated IAM role in the search box and choose the checkbox next to it.
 - i. Choose **Attach policies**.
 - ii. In the search box, enter **AmazonLexRunBotsOnly** and choose the checkbox next to it.

- iii. Choose **Attach policy**.

Next step

[Step 5: Deploy Your Bot as a Web Application](#)

Step 5: Deploy Your Bot as a Web Application

To deploy your bot as a web application



1. Download the repository at https://github.com/awsdocs/amazon-lex-developer-guide/blob/master/example_apps/agent_assistance_bot/ to your computer.
2. Navigate to the downloaded repository and open the `index.html` file in an editor.
3. Make the following changes.
 - a. In the `AWS.config.credentials` section, enter your Region name and your identity pool ID.
 - b. In the `Amazon Lex runtime parameters` section, enter the bot name.
 - c. Save the file.

Step 6: Use the Bot

For demo purposes, you provide input to the bot as the customer and as the agent. To differentiate between the two, questions asked by the customer begin with “Customer:” and answers provided by the agent begin with “Agent:”. You can choose from a menu of suggested inputs.

Run your web application by opening `index.html` to engage in a conversation similar to the following image with your bot:

Call Center Bot with Agent Assistant



The `pushChat()` function in the `index.html` file is explained below.

```
var endConversationStatement = "Customer: I have no more questions. Thank
you."

// If the agent has to send a message, start the message with 'Agent'
var inputText = document.getElementById('input');
if (inputText && inputText.value && inputText.value.trim().length > 0 &&
inputText.value[0]=='Agent') {
    showMessage(inputText.value, 'agentRequest','conversation');
    inputText.value = "";
}
// If the customer has to send a message, start the message with 'Customer'
if(inputText && inputText.value && inputText.value.trim().length > 0 &&
inputText.value[0]=='Customer') {
    // disable input to show we're sending it
    var input = inputText.value.trim();
    inputText.value = '...';
    inputText.locked = true;
    customerInput = input.substring(2);
```

```

        // Send it to the Lex runtime
        var params = {
            botAlias: '$LATEST',
            botName: 'KendraTestBot',
            inputText: customerInput,
            userId: lexUserId,
            sessionAttributes: sessionAttributes
        };

        showMessage(input, 'customerRequest', 'conversation');
        if(input== endConversationStatement){
            showMessage('Conversation
Ended.', 'conversationEndRequest', 'conversation');
        }
        lexruntime.postText(params, function(err, data) {
            if (err) {
                console.log(err, err.stack);
                showMessage('Error:  ' + err.message + ' (see console for
details)', 'lexError', 'conversation1')
            }

            if (data &&input!=endConversationStatement) {
                // capture the sessionAttributes for the next cycle
                sessionAttributes = data.sessionAttributes;

                showMessage(data, 'lexResponse', 'conversation1');
            }
            // re-enable input
            inputText.value = '';
            inputText.locked = false;
        });
    }
    // we always cancel form submission
    return false;

```

When you provide input as a customer, the Amazon Lex runtime API sends it to Amazon Lex.

The `showMessage(daText, senderRequest, displayWindow)` function displays the conversation between the agent and the customer in the chat window. Responses suggested by Amazon Kendra are shown in an adjacent window. The conversation ends when customer says **“I have no more questions. Thank you.”**

Note: Please delete your Amazon Kendra index when not in use.

Migrating a bot

The Amazon Lex V2 API uses an updated information architecture that enables simplified resource versioning and support for multiple languages in a bot. For more information, see the [Migration guide](#) in the *Amazon Lex V2 Developer Guide*.

To use these new features, you need to migrate your bot. When you migrate a bot, Amazon Lex provides the following:

- Migration copies your custom intents and slot types to the Amazon Lex V2 bot.
- You can add multiple languages to the same Amazon Lex V2 bot. In Amazon Lex V1 you create a separate bot for each language. You can migrate multiple Amazon Lex V1 bots, each using a different language, to one Amazon Lex V2 bot.
- Amazon Lex maps Amazon Lex V1 built-in slot types and intents to Amazon Lex V2 built-in slot types and intents. If a built-in can't be migrated, Amazon Lex returns a message that tells you what to do next.

The migration process doesn't migrate the following:

- Aliases
- Amazon Kendra indexes
- AWS Lambda functions
- Conversation log settings
- Messaging channels such as Slack
- Tags

To migrate a bot, your user or role must have IAM permission for both Amazon Lex and Amazon Lex V2 API operations. For the required permissions, see [Allow a user to migrate a bot to Amazon Lex V2 APIs](#).

Migrating a bot (Console)

Use the Amazon Lex V1 console to migrate the structure of a bot to an Amazon Lex V2 bot.

To use the console to migrate a bot to the Amazon Lex V2 API

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. From the left menu, choose **Migration tool**.
3. From the list of bots, choose the bot that you want to migrate and then choose **Migrate**.
4. Choose the version of the bot that you want to migrate, then enter the name of the bot to migrate to. If you enter the name of an existing Amazon Lex V2 bot, the Amazon Lex V1 bot is migrated to the language shown in the details and overwrites the Draft version of the language.
5. Choose **Next**.
6. Choose the IAM role that Amazon Lex uses to run the Amazon Lex V2 API version of the bot. You can choose to create a new role with the minimum permissions required to run the bot, or you can choose an existing IAM role.
7. Choose **Next**.
8. Review the settings for migration. If they look OK, choose **Start migration**.

After you start the migration process, you are returned to the migration tool start page. You can monitor the progress of the migration in the **History** table. When the **Migration status** column says **Complete** the migration is finished.

Amazon Lex uses the `StartImport` operation in the Amazon Lex V2 API to import the migrated bot. You'll see an entry in the Amazon Lex V2 console import history table for each migration.

During the migration, Amazon Lex may find resources in the bot that can't be migrated. You get an error or warning message for each resource that can't be migrated. Each message includes a link to documentation that explains how to resolve the issue.

Migrating a Lambda function

Amazon Lex V2 changes the way that Lambda functions are defined for a bot. It only allows one Lambda function in an alias for each language in a bot. For more information on migrating your Lambda functions, see [Migrating a Lambda function from Amazon Lex V1 to Amazon Lex V2](#).

Migration messages

During migration, Amazon Lex may find resources, such as built-in slot types, that it can't migrate to the equivalent Amazon Lex V2 resource. When this happens, Amazon Lex returns a migration message that describes what happened and provides a link to the documentation that tells you how to fix the migration issue. The following sections describe the issues that might arise when you are migrating a bot and how to fix the issue.

Topics

- [Built-in intent](#)
- [Built-in slot type](#)
- [Conversation logs](#)
- [Message groups](#)
- [Prompts and phrases](#)
- [Other Amazon Lex V1 features](#)

Built-in intent

When you use a built-in intent that is not supported in Amazon Lex V2, the intent is mapped to a custom intent in your Amazon Lex V2 bot. The custom intent doesn't contain utterances. To continue using the intent, add sample utterances.

Built-in slot type

Any migrated slot that uses a slot type that is not supported in Amazon Lex V2 won't be given a slot type value. To use this slot:

- Create a custom slot type
- Add slot type values that are expected for the slot type
- Update the slot to use the new custom slot type

Conversation logs

Migration doesn't update the conversation log settings of the Amazon Lex V2 bot.

To configure conversation logs

1. Open the Amazon Lex V2 console at <https://console.aws.amazon.com/lexv2>.
2. From the list of bots, choose the bot whose conversation logs you want to configure.
3. From the left menu, choose **Aliases**, and then choose an alias from the list.
4. In the **Conversation logs** section, choose **Manage conversation logs** to configure conversation logs for the bot alias.

Message groups

Amazon Lex V2 supports only one message and two alternative messages per message group. If you have more than three messages per message group in an Amazon Lex V1 bot, only the first three messages are migrated. To use more messages in a message group, use a Lambda function to output various messages.

Prompts and phrases

Amazon Lex V2 uses a different mechanism for follow up, clarification, and hang up prompts.

For follow up prompts, use context carryover to switch to a different intent after fulfillment.

For example, suppose that you have an intent to book a car rental that is configured to return an output context called `book_car_fulfilled`. When the intent is fulfilled, Amazon Lex sets the output context variable to `book_car_fulfilled`. Since `book_car_fulfilled` is an active context, an intent with `book_car_fulfilled` as an input context is considered for recognition, as long as the user utterance is recognized as an attempt to elicit that intent. You can use this for intents that only make sense after booking a car, such as emailing a receipt or modifying a reservation.

Amazon Lex V2 does not support clarification prompts and hang up phrases (abort statements). Amazon Lex V2 bots contain a default fallback intent that is invoked if no intents are matched. To send a clarification prompt with retries, configure a Lambda function and enable the dialog code hook in the fallback intent. The Lambda function can output a clarification prompt as a response and the retry value in a session attribute. If the retry value exceeds the maximum number of retries, you can output a hang up phrase and close the conversation.

Other Amazon Lex V1 features

The migration tool supports only migration of Amazon Lex V1 bots and their underlying intents, slot types, and slots. For other features, see the following topics in the Amazon Lex V2 documentation.

- Bot aliases: [Aliases](#)
- Bot channels: [Deploying an Amazon Lex V2 bot on a messaging platform](#)
- Conversation log settings: [Monitoring with conversation logs](#)
- Amazon Kendra indexes: [AMAZON.KendraSearchIntent](#)
- Lambda functions: [Using an AWS Lambda function](#)
- Tags: [Tagging resources](#)

Migrating a Lambda function from Amazon Lex V1 to Amazon Lex V2

Amazon Lex V2 allows only one Lambda function for each language in a bot. The Lambda function and its settings are configured for the bot alias that you use at runtime.

The Lambda function is invoked for all intents in that language if dialog and fulfillment code hooks are enabled for the intent.

Amazon Lex V2 Lambda functions have a different input and output message format from Amazon Lex V1. These are the differences in the Lambda function input format.

- Amazon Lex V2 replaces the `currentIntent` and `alternativeIntents` structures with the `interpretations` structure. Each interpretation contains an intent, the NLU confidence score for the intent, and an optional sentiment analysis.
- Amazon Lex V2 moves the `activeContexts`, `sessionAttributes` in Amazon Lex V1 to the unified `sessionState` structure. This structure provides information about the current state of the conversation, including the originating request ID.
- Amazon Lex V2 doesn't return the `recentIntentSummaryView`. Use the information in the `sessionState` structure instead.
- The Amazon Lex V2 input provides the `botId` and `localeId` in the bot attribute.

- The input structure contains an `inputMode` attribute that provides information on the type of input: text, speech, or DTMF.

These are the differences in the Lambda function output format:

- The `activeContexts` and `sessionAttributes` structures in Amazon Lex V1 are replaced by the `sessionState` structure in Amazon Lex V2.
- The `recentIntentSummaryView` isn't included in the output.
- The Amazon Lex V1 `dialogAction` structure is split into two structures, `dialogAction` that is part of the `sessionState` structure, and `messages` that is required when the `dialogAction.type` is `ElicitIntent`. Amazon Lex chooses messages from this structure to show to the user.

When you build a bot with the Amazon Lex V2 APIs, there is only one Lambda function per bot alias per language instead of a Lambda function for each intent. If you want to continue to use separate functions, you can create a router function that activates a separate function for each intent. The following is a router function that you can use or modify for your application.

```
import os
import json
import boto3

# reuse client connection as global
client = boto3.client('lambda')

def router(event):
    intent_name = event['sessionState']['intent']['name']
    fn_name = os.environ.get(intent_name)
    print(f"Intent: {intent_name} -> Lambda: {fn_name}")
    if (fn_name):
        # invoke lambda and return result
        invoke_response = client.invoke(FunctionName=fn_name, Payload =
json.dumps(event))
        print(invoke_response)
        payload = json.load(invoke_response['Payload'])
        return payload
    raise Exception('No environment variable for intent: ' + intent_name)

def lambda_handler(event, context):
```

```
print(event)
response = router(event)
return response
```

List of updated fields

The following tables provide detailed information about the updated fields in the Amazon Lex V2 Lambda request and response. You can use these tables to map fields between the versions.

Request

The following fields have been updated in the Lambda function request format.

Active contexts

The `activeContexts` structure is now part of the `sessionState` structure.

V1 structure	V2 structure
<code>activeContexts</code>	<code>sessionState.activeContexts</code>
<code>activeContexts[*].timeToLive</code>	<code>sessionState.activeContexts[*].timeToLive</code>
<code>activeContexts[*].timeToLive.timeToLiveInSeconds</code>	<code>sessionState.activeContexts[*].timeToLive.timeToLiveInSeconds</code>
<code>activeContexts[*].timeToLive.turnsToLive</code>	<code>sessionState.activeContexts[*].timeToLive.turnsToLive</code>
<code>activeContexts[*].name</code>	<code>sessionState.activeContexts[*].name</code>
<code>activeContexts[*].parameters</code>	<code>sessionState.activeContexts[*].contextAttributes</code>

Alternative intents

The `interpretations` list from index 1 to N contains the list of alternative intents predicted by Amazon Lex V2, along with their confidence scores. The `recentIntentSummaryView` is removed from the request structure in Amazon Lex V2. To see the details from the `recentIntentSummaryView`, use the [GetSession](#) operation.

V1 structure	V2 structure
alternativeIntents	interpretations[1:*
recentIntentSummaryView	N/A

Bot

In Amazon Lex V2, bots and aliases have identifiers. The bot ID is part of the codehook input. The alias ID is included, but not the alias name. Amazon Lex V2 supports multiple locales for the same bot so the locale ID is included.

V1 structure	V2 structure
bot	bot
bot.name	bot.name
N/A	bot.id
bot.alias	N/A
N/A	bot.aliasId
bot.version	bot.version
N/A	bot.localeId

Current intent

The `sessionState.intent` structure contains the details of the active intent. Amazon Lex V2 also returns a list of all of the intents, including alternative intents, in the `interpretations` structure. The first element in the `interpretations` list is always the same as `sessionState.intent`.

V1 structure	V2 structure
currentIntent	sessionState.intent OR interpretations[0].intent
currentIntent.name	sessionState.intent.name OR interpretations[0].intent.name
currentIntent.nluConfidenceScore	interpretations[0].nluConfidence.score

Dialog action

The `confirmationStatus` field is now part of the `sessionState` structure.

V1 structure	V2 structure
currentIntent.confirmationStatus	sessionState.intent.confirmationState OR interpretations[0].intent.confirmationState
N/A	sessionState.intent.state OR interpretations[*].intent.state

Amazon Kendra

The `kendraResponse` field is now part of the `sessionState` and `interpretations` structures.

V1 structure	V2 structure
kendraResponse	sessionState.intent.kendraResponse OR interpretations[0].intent.kendraResponse

Sentiment

The `sentimentResponse` structure is moved to the new `interpretations` structure.

V1 structure	V2 structure
sentimentResponse	interpretations[0].sentimentResponse
sentimentResponse.sentimentLabel	interpretations[0].sentimentResponse .sentiment
sentimentResponse.sentimentScore	interpretations[0].sentimentResponse .sentimentScore

Slots

Amazon Lex V2 provides a single `slots` object inside the `sessionState.intent` structure that contains the resolved values, interpreted value, and the original value of what the user said. Amazon Lex V2 also supports multi-valued slots by setting the `slotShape` as `List` and setting the `values` list. Single-value slots are supported by the `value` field, their shape is assumed to be `Scalar`.

V1 structure	V2 structure
currentIntent.slots	sessionState.intent.slots OR interpretations[0].intent.slots
currentIntent.slots[*].value	sessionState.intent.slots[*].value.interpretedValue OR interpretations[0].intent.slots[*].value.interpretedValue
N/A	sessionState.intent.slots[*].value.shape OR interpretations[0].intent.slots[*].shape
N/A	sessionState.intent.slots[*].values OR interpretations[0].intent.slots[*].values
currentIntent.slotDetails	sessionState.intent.slots OR interpretations[0].intent.slots

V1 structure	V2 structure
currentIntent.slotDetails[*].resolutions	sessionState.intent.slots[*].resolvedValues OR interpretations[0].intent.slots[*].resolvedValues
currentIntent.slotDetails[*].originalValue	sessionState.intent.slots[*].originalValue OR interpretations[0].intent.slots[*].originalValue

Others

The Amazon Lex V2 `sessionId` field is the same as the `userId` field in Amazon Lex V1. Amazon Lex V2 also sends the `inputMode` of the caller: text, DTMF, or speech.

V1 structure	V2 structure
userId	sessionId
inputTranscript	inputTranscript
invocationSource	invocationSource
outputDialogMode	responseContentType
messageVersion	messageVersion
sessionAttributes	sessionState.sessionAttributes
requestAttributes	requestAttributes
N/A	inputMode
N/A	originatingRequestId

Response

The following fields have been changed in the Lambda function response message format.

Active contexts

The `activeContexts` structure moved to the `sessionState` structure.

V1 structure	V2 structure
<code>activeContexts</code>	<code>sessionState.activeContexts</code>
<code>activeContexts[*].timeToLive</code>	<code>sessionState.activeContexts[*].timeToLive</code>
<code>activeContexts[*].timeToLive.timeToLiveInSeconds</code>	<code>sessionState.activeContexts[*].timeToLive.timeToLiveInSeconds</code>
<code>activeContexts[*].timeToLive.turnsToLive</code>	<code>sessionState.activeContexts[*].timeToLive.turnsToLive</code>
<code>activeContexts[*].name</code>	<code>sessionState.activeContexts[*].name</code>
<code>activeContexts[*].parameters</code>	<code>sessionState.activeContexts[*].contextAttributes</code>

Dialog action

The `dialogAction` structure moved to the `sessionState` structure. You can now specify multiple messages in a dialog action, and the `genericAttachments` structure is now the `imageResponseCard` structure.

V1 structure	V2 structure
<code>dialogAction</code>	<code>sessionState.dialogAction</code>
<code>dialogAction.type</code>	<code>sessionState.dialogAction.type</code>
<code>dialogAction.slotToElicit</code>	<code>sessionState.intent.dialogAction.slotToElicit</code>
<code>dialogAction.type.fulfillmentState</code>	<code>sessionState.intent.state</code>
<code>dialogAction.message</code>	<code>messages</code>
<code>dialogAction.message.contentType</code>	<code>messages[*].contentType</code>

V1 structure	V2 structure
dialogAction.message.content	messages[*].content
dialogAction.responseCard	messages[*].imageResponseCard
dialogAction.responseCard.version	N/A
dialogAction.responseCard.contentType	messages[*].contentType
dialogAction.responseCard.genericAttachments	N/A
dialogAction.responseCard.genericAttachments[*].title	messages[*].imageResponseCard.title
dialogAction.responseCard.genericAttachments[*].subTitle	messages[*].imageResponseCard.subtitle
dialogAction.responseCard.genericAttachments[*].imageUrl	messages[*].imageResponseCard.imageUrl
dialogAction.responseCard.genericAttachments[*].buttons	messages[*].imageResponseCard.buttons
dialogAction.responseCard.genericAttachments[*].buttons[*].value	messages[*].imageResponseCard.buttons[*].value
dialogAction.responseCard.genericAttachments[*].buttons[*].text	messages[*].imageResponseCard.buttons[*].text
dialogAction.kendraQueryRequestPayload	dialogAction.kendraQueryRequestPayload
dialogAction.kendraQueryFilterString	dialogAction.kendraQueryFilterString

Intents and slots

Intent and slot fields that were part of the `dialogAction` structure are now part of the `sessionState` structure.

V1 structure	V2 structure
dialogAction.intentName	sessionState.intent.name
dialogAction.slots	sessionState.intent.slots
dialogAction.slots[*].key	sessionState.intent.slots[*].key
dialogAction.slots[*].value	sessionState.intent.slots[*].value.interpretedValue
N/A	sessionState.intent.slots[*].value.shape
N/A	sessionState.intent.slots[*].values

Others

The `sessionAttributes` structure is now part of the `sessionState` structure. The `recentIntentSummaryReview` structure has been removed.

V1 structure	V2 structure
sessionAttributes	sessionState.sessionAttributes
recentIntentSummaryView	N/A

Security in Amazon Lex

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. The effectiveness of our security is regularly tested and verified by third-party auditors as part of the [AWS compliance programs](#). To learn about the compliance programs that apply to Amazon Lex, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your organization's requirements, and applicable laws and regulations.

This documentation will help you understand how to apply the shared responsibility model when using Amazon Lex. The following topics show you how to configure Amazon Lex to meet your security and compliance objectives. You'll also learn how to use other AWS services that can help you to monitor and secure your Amazon Lex resources.

Topics

- [Data Protection in Amazon Lex](#)
- [Identity and Access Management for Amazon Lex](#)
- [Monitoring in Amazon Lex](#)
- [Compliance Validation for Amazon Lex](#)
- [Resilience in Amazon Lex](#)
- [Infrastructure Security in Amazon Lex](#)

Data Protection in Amazon Lex

Amazon Lex collects customer content for troubleshooting and to help improve the service. Customer content is secured by default. You can delete content for individual customers using the Amazon Lex API.

Amazon Lex stores four types of content:

- Sample utterances, which are used to build and train a bot
- Customer utterances from users interacting with the bot
- Session attributes, which provide application-specific information for the duration of a user's interaction with a bot
- Request attributes, which contain information that applies to a single request to a bot

Any Amazon Lex bot that is designed for use by children is governed by the Children's Online Privacy Protection Act (COPPA). You tell Amazon Lex that the bot is subject to COPPA by using the console or the Amazon Lex API to set the `childDirected` field to `true`. When the `childDirected` field is set to `true`, no user utterances are stored.

Topics

- [Encryption at Rest](#)
- [Encryption in Transit](#)
- [Key Management](#)

Encryption at Rest

Amazon Lex encrypts the user utterances that it stores.

Topics

- [Sample Utterances](#)
- [Customer Utterances](#)
- [Session Attributes](#)
- [Request Attributes](#)

Sample Utterances

When you develop a bot, you can provide sample utterances for each intent and slot. You can also provide custom values and synonyms for slots. This information is encrypted at rest, and it is used to build the bot and to create the user experience.

Customer Utterances

Amazon Lex encrypts utterances that users send to your bot unless the `childDirected` field is set to `true`.

When the `childDirected` field is set to `true`, no user utterances are stored.

When the `childDirected` field is set to `false` (the default), user utterances are encrypted and stored for 15 days for use with the [GetUtterancesView](#) operation. To delete stored utterances for a specific user, use the [DeleteUtterances](#) operation .

When your bot accepts voice input, the input is stored indefinitely. Amazon Lex uses it to improve your bot's ability to respond to user input.

Use the [DeleteUtterances](#) operation to delete stored utterances for a specific user.

Session Attributes

Session attributes contain application-specific information that is passed between Amazon Lex and client applications. Amazon Lex passes session attributes to all AWS Lambda functions configured for a bot. If a Lambda function adds or updates session attributes, Amazon Lex passes the new information back to the client application.

Session attributes persist in an encrypted store for the duration of the session. You can configure the session to remain active for a minimum of 1 minute and up to 24 hours after the last user utterance. The default session duration is 5 minutes.

Request Attributes

Request attributes contain request-specific information and apply only to the current request. A client application uses request attributes to send information to Amazon Lex at runtime.

You use request attributes to pass information that doesn't need to persist for the entire session. Because request attributes don't persist across requests, they aren't stored.

Encryption in Transit

Amazon Lex uses the HTTPS protocol to communicate with your client application. It uses HTTPS and AWS signatures to communicate with other services, such as Amazon Polly and AWS Lambda on your application's behalf.

Key Management

Amazon Lex protects your content from unauthorized use with internal keys.

Identity and Access Management for Amazon Lex

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Lex resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How Amazon Lex works with IAM](#)
- [Identity-based policy examples for Amazon Lex](#)
- [AWS managed policies for Amazon Lex](#)
- [Using Service-Linked Roles for Amazon Lex](#)
- [Troubleshooting Amazon Lex identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Amazon Lex.

Service user – If you use the Amazon Lex service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon Lex features to do your work, you might need additional permissions. Understanding how access is managed can

help you request the right permissions from your administrator. If you cannot access a feature in Amazon Lex, see [Troubleshooting Amazon Lex identity and access](#).

Service administrator – If you're in charge of Amazon Lex resources at your company, you probably have full access to Amazon Lex. It's your job to determine which Amazon Lex features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon Lex, see [How Amazon Lex works with IAM](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon Lex. To view example Amazon Lex identity-based policies that you can use in IAM, see [Identity-based policy examples for Amazon Lex](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see [AWS Signature Version 4 for API requests](#) in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Multi-factor authentication](#) in

the *AWS IAM Identity Center User Guide* and [AWS Multi-factor authentication in IAM](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see [What is IAM Identity Center?](#) in the *AWS IAM Identity Center User Guide*.

IAM users and groups

An [IAM user](#) is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An [IAM group](#) is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier

to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the AWS Management Console, you can [switch from a user to an IAM role \(console\)](#). You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see [Create a role for a third-party identity provider \(federation\)](#) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see [Permission sets](#) in the *AWS IAM Identity Center User Guide*.
- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.

- **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).
- **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Create a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.
- **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Use an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Choose between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see [Service control policies](#) in the *AWS Organizations User Guide*.
- **Resource control policies (RCPs)** – RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the AWS account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of AWS services that support RCPs, see [Resource control policies \(RCPs\)](#) in the *AWS Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's

permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Amazon Lex works with IAM

Before you use IAM to manage access to Amazon Lex, learn what IAM features are available to use with Amazon Lex.

IAM features you can use with Amazon Lex

IAM feature	Amazon Lex support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys (service-specific)	Yes
ACLs	No
ABAC (tags in policies)	Partial
Temporary credentials	Yes
Principal permissions	Yes
Service roles	Yes
Service-linked roles	Yes

To get a high-level view of how Amazon Lex and other AWS services work with most IAM features, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Identity-based policies for Amazon Lex

Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Identity-based policy examples for Amazon Lex

To view examples of Amazon Lex identity-based policies, see [Identity-based policy examples for Amazon Lex](#).

Resource-based policies within Amazon Lex

Supports resource-based policies: No

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource

are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Policy actions for Amazon Lex

Supports policy actions: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Amazon Lex actions, see [Actions defined by Amazon Lex](#) in the *Service Authorization Reference*.

Policy actions in Amazon Lex use the following prefix before the action:

```
lex
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [  
    "lex:action1",  
    "lex:action2"  
]
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word Describe, include the following action:

```
"Action": "lex:Describe"
```

Policy resources for Amazon Lex

Supports policy resources: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": ""
```

An Amazon Lex bot resource ARN has the following format.

```
arn:aws:lex:${Region}:${Account}:bot:${Bot-Name}
```

For more information about the format of ARNs, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#).

For example, to specify the OrderFlowers bot in your statement, use the following ARN.

```
"Resource": "arn:aws:lex:us-east-2:123456789012:bot:OrderFlowers"
```

To specify all bots that belong to a specific account, use the wildcard (*).

```
"Resource": "arn:aws:lex:us-east-2:123456789012:bot:*"
```

Some Amazon Lex actions, such as those for creating resources, can't be performed on a specific resource. In those cases, you must use the wildcard, (*).

```
"Resource": ""
```

To see a list of Amazon Lex resource types and their ARNs, see [Resources defined by Amazon Lex](#) in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see [Actions defined by Amazon Lex](#).

Policy condition keys for Amazon Lex

Supports service-specific policy condition keys: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or *Condition block*) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

To see a list of Amazon Lex condition keys, see [Condition keys for Amazon Lex](#) in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions defined by Amazon Lex](#).

The following table lists the Amazon Lex condition keys that apply to Amazon Lex resources. You can include these keys in `Condition` elements in an IAM permissions policy.

ACLs in Amazon Lex

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

ABAC with Amazon Lex

Supports ABAC (tags in policies): Partial

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [Define permissions with ABAC authorization](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control \(ABAC\)](#) in the *IAM User Guide*.

You can associate tags with certain types of Amazon Lex resources for authorization. To control access based on tags, provide tag information in the condition element of a policy by using the `lex:ResourceTag/${TagKey}`, `aws:RequestTag/${TagKey}`, or `aws:TagKeys` condition keys.

For more information about tagging Amazon Lex resources, see [Tagging Your Amazon Lex Resources](#).

To view an example identity-based policy for limiting access to a resource based on the tags on that resource, see [Use a Tag to Access a Resource](#).

The following table lists the actions and corresponding resource types for tag-based access control. Each action is authorized based on the tags associated with the corresponding resource type.

Action	Resource type	Condition keys	Notes
CreateBotVersion	bot	lex:ResourceTag	
DeleteBot	bot	lex:ResourceTag	
DeleteBotAlias	alias	lex:ResourceTag	
DeleteBotChannelAssociation	channel	lex:ResourceTag	
DeleteBotVersion	bot	lex:ResourceTag	
DeleteSession	bot or alias	lex:ResourceTag	Uses tags associated with the bot when alias is set to \$LATEST. Uses tags associated with the specified alias when used with other aliases.
DeleteUtterances	bot	lex:ResourceTag	
GetBot	bot or alias	lex:ResourceTag	Uses tags associated with the bot when versionOrAlias is set to \$LATEST or numeric version. Uses tags associated with the specified alias when used with aliases
GetBotAlias	alias	lex:ResourceTag	

Action	Resource type	Condition keys	Notes
GetBotChannelAssociation	channel	lex:ResourceTag	
GetBotChannelAssociations	channel	lex:ResourceTag	Uses tags associated with the bot when alias is set to "-". Uses tags associated with the specified alias when a bot alias is specified
GetBotVersions	bot	lex:ResourceTag	
GetExport	bot	lex:ResourceTag	
GetSession	bot or alias	lex:ResourceTag	Uses tags associated with the bot when alias is set to \$LATEST. Uses tags associated with the specified alias when used with other aliases.
GetUtterancesView	bot	lex:ResourceTag	
ListTagsForResource	bot, alias, or channel	lex:ResourceTag	
PostContent	bot or alias	lex:ResourceTag	Uses tags associated with the bot when alias is set to \$LATEST. Uses tags associated with the specified alias when used with other aliases.

Action	Resource type	Condition keys	Notes
PostText	bot or alias	lex:ResourceTag	Uses tags associated with the bot when alias is set to \$LATEST. Uses tags associated with the specified alias when used with other aliases.
PutBot	bot	lex:ResourceTag, aws:RequestTag, aws:TagKeys	
PutBotAlias	alias	lex:ResourceTag, aws:RequestTag, aws:TagKeys	
PutSession	bot or alias	lex:ResourceTag	Uses tags associated with the bot when alias is set to \$LATEST. Uses tags associated with the specified alias when used with other aliases.
StartImport	bot	lex:ResourceTag	Relies on access policy for the PutBot operation. Tags and permissions specific to the StartImport operation are ignored.

Action	Resource type	Condition keys	Notes
TagResource	bot, alias, or channel	lex:ResourceTag, aws:RequestTag, aws:TagKeys	
UntagResource	bot, alias, or channel	lex:ResourceTag, aws:RequestTag, aws:TagKeys	

Using temporary credentials with Amazon Lex

Supports temporary credentials: Yes

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see [AWS services that work with IAM](#) in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see [Switch from a user to an IAM role \(console\)](#) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#).

You can use temporary credentials to sign in with federation, assume an IAM role, or to assume a cross-account role. You obtain temporary security credentials by calling AWS STS API operations such as [AssumeRole](#) or [GetFederationToken](#).

Cross-service principal permissions for Amazon Lex

Supports forward access sessions (FAS): Yes

When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).

Service roles for Amazon Lex

Supports service roles: Yes

A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Create a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

Warning

Changing the permissions for a service role might break Amazon Lex functionality. Edit service roles only when Amazon Lex provides guidance to do so.

Choosing an IAM role in Amazon Lex

Amazon Lex uses service-linked roles to call Amazon Comprehend and Amazon Polly. It uses resource-level permissions on your AWS Lambda functions to invoke them.

You must provide an IAM role to enable conversation tagging. For more information, see [Creating an IAM Role and Policies for Conversation Logs](#).

Service-linked roles for Amazon Lex

Supports service-linked roles: Yes

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing Amazon Lex service-linked roles, see [Using Service-Linked Roles for Amazon Lex](#).

Identity-based policy examples for Amazon Lex

By default, users and roles don't have permission to create or modify Amazon Lex resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see [Create IAM policies \(console\)](#) in the *IAM User Guide*.

For details about actions and resource types defined by Amazon Lex, including the format of the ARNs for each of the resource types, see [Actions, resources, and condition keys for Amazon Lex](#) in the *Service Authorization Reference*.

Topics

- [Policy best practices](#)
- [Using the Amazon Lex console](#)
- [Allow users to view their own permissions](#)
- [Delete All Amazon Lex Bots](#)
- [Allow a user to migrate a bot to Amazon Lex V2 APIs](#)
- [Use a Tag to Access a Resource](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon Lex resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.
- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on

specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.

- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see [Validate policies with IAM Access Analyzer](#) in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Secure API access with MFA](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Using the Amazon Lex console

To access the Amazon Lex console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon Lex resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

AWS addresses many common use cases by providing standalone IAM policies that are created and administered by AWS. These policies are called AWS managed policies. AWS managed policies make it easier for you to assign appropriate permissions to users, groups, and roles than if you had

to write the policies yourself. For more information, see [AWS Managed Policies](#) in the *IAM User Guide*.

The following AWS managed policies, which you can attach to groups and roles in your account, are specific to Amazon Lex:

- **AmazonLexReadOnly** — Grants read-only access to Amazon Lex resources.
- **AmazonLexRunBotsOnly** — Grants access to run Amazon Lex conversational bots.
- **AmazonLexFullAccess** — Grants full access to create, read, update, delete, and run all Amazon Lex resources. Also grants the ability to associate Lambda functions whose name starts with AmazonLex with Amazon Lex intents.

Note

You can review these permissions policies by signing in to the IAM console and searching for specific policies.

The **AmazonLexFullAccess** policy doesn't grant the user permission to use the `KendraSearchIntent` intent to query an Amazon Kendra index. To query an index, you must add additional permissions to the policy. For the required permissions, see [IAM Policy for Amazon Kendra Search](#).

You can also create your own custom IAM policies to allow permissions for Amazon Lex API actions. You can attach these custom policies to the IAM roles or groups that require those permission.

For details about AWS managed policies for Amazon Lex, see [AWS managed policies for Amazon Lex](#).

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```

        "Sid": "ViewOwnUserInfo",
        "Effect": "Allow",
        "Action": [
            "iam:GetUserPolicy",
            "iam:ListGroupsForUser",
            "iam:ListAttachedUserPolicies",
            "iam:ListUserPolicies",
            "iam:GetUser"
        ],
        "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    },
    {
        "Sid": "NavigateInConsole",
        "Effect": "Allow",
        "Action": [
            "iam:GetGroupPolicy",
            "iam:GetPolicyVersion",
            "iam:GetPolicy",
            "iam:ListAttachedGroupPolicies",
            "iam:ListGroupPolicies",
            "iam:ListPolicyVersions",
            "iam:ListPolicies",
            "iam:ListUsers"
        ],
        "Resource": "*"
    }
]
}

```

Delete All Amazon Lex Bots

This example policy grants a user in your AWS account permission to delete any bot in your account.

```

{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "lex:DeleteBot"
            ],
            "Resource": [

```



```

        "*"
    ]
}
]
}

```

Allow a user to migrate a bot to Amazon Lex V2 APIs

The following IAM permission policy allows a user to start migrating a bot from Amazon Lex to Amazon Lex V2 APIs and to see the list of migrations and their progress.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "startMigration",
      "Effect": "Allow",
      "Action": "lex:StartMigration",
      "Resource": "arn:aws:lex:<Region>:<123456789012>:bot:*"
    },
    {
      "Sid": "passRole",
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::<123456789012>:role/<v2 bot role>"
    },
    {
      "Sid": "allowOperations",
      "Effect": "Allow",
      "Action": [
        "lex:CreateBot",
        "lex:CreateIntent",
        "lex:UpdateSlot",
        "lex:DescribeBotLocale",
        "lex:UpdateBotAlias",
        "lex:CreateSlotType",
        "lex>DeleteBotLocale",
        "lex:DescribeBot",
        "lex:UpdateBotLocale",
        "lex:CreateSlot",
        "lex>DeleteSlot",
        "lex:UpdateBot",
        "lex>DeleteSlotType",

```

```

        "lex:DescribeBotAlias",
        "lex:CreateBotLocale",
        "lex>DeleteIntent",
        "lex:StartImport",
        "lex:UpdateSlotType",
        "lex:UpdateIntent",
        "lex:DescribeImport",
        "lex:CreateCustomVocabulary",
        "lex:UpdateCustomVocabulary",
        "lex>DeleteCustomVocabulary",
        "lex:DescribeCustomVocabulary",
        "lex:DescribeCustomVocabularyMetadata"
    ],
    "Resource": [
        "arn:aws:lex:<Region>:<123456789012>:bot/*",
        "arn:aws:lex:<Region>:<123456789012>:bot-alias/*/*"
    ]
},
{
    "Sid": "showBots",
    "Effect": "Allow",
    "Action": [
        "lex:CreateUploadUrl",
        "lex:ListBots"
    ],
    "Resource": "*"
},
{
    "Sid": "showMigrations",
    "Effect": "Allow",
    "Action": [
        "lex:GetMigration",
        "lex:GetMigrations"
    ],
    "Resource": "*"
}
]
}
```

Use a Tag to Access a Resource

This example policy grants a user or role in your AWS account permission to use the `PostText` operation with any resource tagged with the key **Department** and the value **Support**.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "lex:PostText",
      "Effect": "Allow",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "lex:ResourceTag/Department": "Support"
        }
      }
    }
  ]
}
```

AWS managed policies for Amazon Lex

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

AWS managed policy: AmazonLexReadOnly

You can attach the AmazonLexReadOnly policy to your IAM identities.

This policy grants read-only permissions that allow users to view all actions in the Amazon Lex and Amazon Lex V2 model building service.

Permissions details

This policy includes the following permissions:

- `lex` – Read-only access to Amazon Lex and Amazon Lex V2 resources in the model building service.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "lex:GetBot",
        "lex:GetBotAlias",
        "lex:GetBotAliases",
        "lex:GetBots",
        "lex:GetBotChannelAssociation",
        "lex:GetBotChannelAssociations",
        "lex:GetBotVersions",
        "lex:GetBuiltinIntent",
        "lex:GetBuiltinIntents",
        "lex:GetBuiltinSlotTypes",
        "lex:GetIntent",
        "lex:GetIntents",
        "lex:GetIntentVersions",
        "lex:GetSlotType",
        "lex:GetSlotTypes",
        "lex:GetSlotTypeVersions",
        "lex:GetUtterancesView",
        "lex:DescribeBot",
        "lex:DescribeBotAlias",
        "lex:DescribeBotChannel",
        "lex:DescribeBotLocale",
        "lex:DescribeBotVersion",
        "lex:DescribeExport",
```

```

        "lex:DescribeImport",
        "lex:DescribeIntent",
        "lex:DescribeResourcePolicy",
        "lex:DescribeSlot",
        "lex:DescribeSlotType",
        "lex:ListBots",
        "lex:ListBotLocales",
        "lex:ListBotAliases",
        "lex:ListBotChannels",
        "lex:ListBotVersions",
        "lex:ListBuiltInIntents",
        "lex:ListBuiltInSlotTypes",
        "lex:ListExports",
        "lex:ListImports",
        "lex:ListIntents",
        "lex:ListSlots",
        "lex:ListSlotTypes",
        "lex:ListTagsForResource"
    ],
    "Resource": "*"
}
]
}

```

AWS managed policy: AmazonLexRunBotsOnly

You can attach the AmazonLexRunBotsOnly policy to your IAM identities.

This policy grants read-only permissions that allow access to run Amazon Lex and Amazon Lex V2 conversational bots.

Permissions details

This policy includes the following permissions:

- `lex` – Read-only access to all actions in the Amazon Lex and Amazon Lex V2 runtime.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [

```

```

        "lex:PostContent",
        "lex:PostText",
        "lex:PutSession",
        "lex:GetSession",
        "lex>DeleteSession",
        "lex:RecognizeText",
        "lex:RecognizeUtterance",
        "lex:StartConversation"
    ],
    "Resource": "*"
}
]
}

```

AWS managed policy: AmazonLexFullAccess

You can attach the `AmazonLexFullAccess` policy to your IAM identities.

This policy grants administrative permissions that allow the user permission to create, read, update, and delete Amazon Lex and Amazon Lex V2 resources, and to run Amazon Lex and Amazon Lex V2 conversational bots.

Permissions details

This policy includes the following permissions:

- `lex` – Allows principals read and write access to all actions in the Amazon Lex and Amazon Lex V2 model building and runtime services.
- `cloudwatch` – Allows principals to view Amazon CloudWatch metrics and alarms.
- `iam` – Allows principals to create and delete service-linked roles, pass roles, and attach and detach policies to a role. The permissions are restricted to `"lex.amazonaws.com"` for Amazon Lex operations and to `"lexv2.amazonaws.com"` for Amazon Lex V2 operations.
- `kendra` – Allows principals to list Amazon Kendra indexes.
- `kms` – Allows principals to describe AWS KMS keys and aliases.
- `lambda` – Allows principals to list AWS Lambda functions and manage permissions attached to any Lambda function.
- `polly` – Allows principals to describe Amazon Polly voices and synthesize speech.

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "cloudwatch:GetMetricStatistics",
      "cloudwatch:DescribeAlarms",
      "cloudwatch:DescribeAlarmsForMetric",
      "kms:DescribeKey",
      "kms:ListAliases",
      "lambda:GetPolicy",
      "lambda:ListFunctions",
      "lex:*",
      "polly:DescribeVoices",
      "polly:SynthesizeSpeech",
      "kendra:ListIndices",
      "iam:ListRoles",
      "s3:ListAllMyBuckets",
      "logs:DescribeLogGroups",
      "s3:GetBucketLocation"
    ],
    "Resource": [
      "*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "lambda:AddPermission",
      "lambda:RemovePermission"
    ],
    "Resource": "arn:aws:lambda:*:*:function:AmazonLex*",
    "Condition": {
      "StringEquals": {
        "lambda:Principal": "lex.amazonaws.com"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": [
      "iam:GetRole"
    ],
    "Resource": [

```

```

        "arn:aws:iam::*:role/aws-service-role/lex.amazonaws.com/
AWSServiceRoleForLexBots",
        "arn:aws:iam::*:role/aws-service-role/channels.lex.amazonaws.com/
AWSServiceRoleForLexChannels",
        "arn:aws:iam::*:role/aws-service-role/lexv2.amazonaws.com/
AWSServiceRoleForLexV2Bots*",
        "arn:aws:iam::*:role/aws-service-role/channels.lexv2.amazonaws.com/
AWSServiceRoleForLexV2Channels*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "iam:CreateServiceLinkedRole"
    ],
    "Resource": [
        "arn:aws:iam::*:role/aws-service-role/lex.amazonaws.com/
AWSServiceRoleForLexBots"
    ],
    "Condition": {
        "StringEquals": {
            "iam:AWSServiceName": "lex.amazonaws.com"
        }
    }
},
{
    "Effect": "Allow",
    "Action": [
        "iam:CreateServiceLinkedRole"
    ],
    "Resource": [
        "arn:aws:iam::*:role/aws-service-role/channels.lex.amazonaws.com/
AWSServiceRoleForLexChannels"
    ],
    "Condition": {
        "StringEquals": {
            "iam:AWSServiceName": "channels.lex.amazonaws.com"
        }
    }
},
{
    "Effect": "Allow",
    "Action": [
        "iam:CreateServiceLinkedRole"
    ]
}

```



```

    ],
    "Resource": [
        "arn:aws:iam::*:role/aws-service-role/lexv2.amazonaws.com/
AWSServiceRoleForLexV2Bots*"
    ],
    "Condition": {
        "StringEquals": {
            "iam:AWSServiceName": "lexv2.amazonaws.com"
        }
    }
},
{
    "Effect": "Allow",
    "Action": [
        "iam:CreateServiceLinkedRole"
    ],
    "Resource": [
        "arn:aws:iam::*:role/aws-service-role/channels.lexv2.amazonaws.com/
AWSServiceRoleForLexV2Channels*"
    ],
    "Condition": {
        "StringEquals": {
            "iam:AWSServiceName": "channels.lexv2.amazonaws.com"
        }
    }
},
{
    "Effect": "Allow",
    "Action": [
        "iam>DeleteServiceLinkedRole",
        "iam:GetServiceLinkedRoleDeletionStatus"
    ],
    "Resource": [
        "arn:aws:iam::*:role/aws-service-role/lex.amazonaws.com/
AWSServiceRoleForLexBots",
        "arn:aws:iam::*:role/aws-service-role/channels.lex.amazonaws.com/
AWSServiceRoleForLexChannels",
        "arn:aws:iam::*:role/aws-service-role/lexv2.amazonaws.com/
AWSServiceRoleForLexV2Bots*",
        "arn:aws:iam::*:role/aws-service-role/channels.lexv2.amazonaws.com/
AWSServiceRoleForLexV2Channels*"
    ]
},
{

```

```

        "Effect": "Allow",
        "Action": [
            "iam:PassRole"
        ],
        "Resource": [
            "arn:aws:iam::*:role/aws-service-role/lex.amazonaws.com/
AWSServiceRoleForLexBots"
        ],
        "Condition": {
            "StringEquals": {
                "iam:PassedToService": [
                    "lex.amazonaws.com"
                ]
            }
        }
    },
    {
        "Effect": "Allow",
        "Action": [
            "iam:PassRole"
        ],
        "Resource": [
            "arn:aws:iam::*:role/aws-service-role/lexv2.amazonaws.com/
AWSServiceRoleForLexV2Bots*"
        ],
        "Condition": {
            "StringEquals": {
                "iam:PassedToService": [
                    "lexv2.amazonaws.com"
                ]
            }
        }
    },
    {
        "Effect": "Allow",
        "Action": [
            "iam:PassRole"
        ],
        "Resource": [
            "arn:aws:iam::*:role/aws-service-role/channels.lexv2.amazonaws.com/
AWSServiceRoleForLexV2Channels*"
        ],
        "Condition": {
            "StringEquals": {

```

```

    "iam:PassedToService": [
      "channels.lexv2.amazonaws.com"
    ]
  }
}

```

Amazon Lex updates to AWS managed policies

View details about updates to AWS managed policies for Amazon Lex since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Amazon Lex [Document History for Amazon Lex](#) page.

Change	Description	Date
<u>AmazonLexFullAccess</u> – Update to an existing policy	Amazon Lex added new permissions to allow read-only access to Amazon Lex V2 model building service operations.	August 18, 2021
<u>AmazonLexReadOnly</u> – Update to an existing policy	Amazon Lex added new permissions to allow read-only access to Amazon Lex V2 model building service operations.	August 18, 2021
<u>AmazonLexRunBotsOnly</u> – Update to an existing policy	Amazon Lex added new permissions to allow read-only access to Amazon Lex V2 runtime service operations.	August 18, 2021
Amazon Lex started tracking changes	Amazon Lex started tracking changes for its AWS managed policies.	August 18, 2021

Using Service-Linked Roles for Amazon Lex

Amazon Lex uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to Amazon Lex. Service-linked roles are predefined by Amazon Lex and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up Amazon Lex easier because you don't have to manually add the necessary permissions. Amazon Lex defines the permissions of its service-linked roles, and unless defined otherwise, only Amazon Lex can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting its related resources. This protects your Amazon Lex resources because you can't inadvertently remove permission to access the resources.

Service-Linked Roles Permissions for Amazon Lex

Amazon Lex uses two service linked roles:

- **AWSServiceRoleForLexBots** – Amazon Lex uses this service-linked role to invoke Amazon Polly to synthesize speech responses for your bot, to call Amazon Comprehend for sentiment analysis, and optionally Amazon Kendra for searching indexes.
- **AWSServiceRoleForLexChannels** – Amazon Lex uses this service-linked role to post text to your bot when managing channels.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-Linked Role Permissions](#) in the *IAM User Guide*.

Creating a Service-Linked Role for Amazon Lex

You don't need to manually create a service-linked role. When you create a bot, bot channel, or Amazon Kendra search intent in the AWS Management Console, Amazon Lex creates the service-linked role for you.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create a new bot, channel association, or Amazon Kendra search intent, Amazon Lex creates the service-linked role for you again.

You can also use the AWS CLI to create a service-linked role with the **AWSServiceRoleForLexBots** use case. In the AWS CLI create a service-linked role with the Amazon Lex service name `lex.amazonaws.com`. For more information, see [Step 1: Create a Service-Linked Role \(AWS CLI\)](#). If you delete this service-linked role, you can use this same process to create the role again.

Editing a Service-Linked Role for Amazon Lex

Amazon Lex does not allow you to edit Amazon Lex service-linked roles. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a Service-Linked Role](#) in the *IAM User Guide*.

Deleting a Service-Linked Role for Amazon Lex

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

Note

If the Amazon Lex service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes and try the operation again.

To delete Amazon Lex resources used by service-linked roles:

1. Delete any bot channels that you are using.
2. Delete any bots in your account.

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the Amazon Lex service-linked roles. For more information, see [Deleting a Service-Linked Role](#) in the *IAM User Guide*.

Supported Regions for Amazon Lex Service-Linked Roles

Amazon Lex supports using service-linked roles in all of the regions where the service is available. For more information, see [Amazon Lex endpoints and quotas](#).

Troubleshooting Amazon Lex identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon Lex and IAM.

Topics

- [I am not authorized to perform an action in Amazon Lex](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my AWS account to access my Amazon Lex resources](#)

I am not authorized to perform an action in Amazon Lex

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the mateojackson IAM user tries to use the console to view details about a fictional *my-example-widget* resource but doesn't have the fictional `lex:GetWidget` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
lex:GetWidget on resource: my-example-widget
```

In this case, the policy for the mateojackson user must be updated to allow access to the *my-example-widget* resource by using the `lex:GetWidget` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon Lex.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Amazon Lex. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Amazon Lex resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon Lex supports these features, see [How Amazon Lex works with IAM](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Monitoring in Amazon Lex

Monitoring is important for maintaining the reliability, availability, and performance of your Amazon Lex chatbots. This topic describes how to use Amazon CloudWatch Logs and AWS CloudTrail to monitor Amazon Lex and describes the Amazon Lex runtime and channel association metrics.

Topics

- [Monitoring Amazon Lex with Amazon CloudWatch](#)
- [Monitoring Amazon Lex API Calls with AWS CloudTrail Logs](#)

Monitoring Amazon Lex with Amazon CloudWatch

To track the health of your Amazon Lex bots, use Amazon CloudWatch. With CloudWatch, you can get metrics for individual Amazon Lex operations or for global Amazon Lex operations for your account. You can also set up CloudWatch alarms to be notified when one or more metrics exceeds a threshold that you define. For example, you can monitor the number of requests made to a bot over a particular time period, view the latency of successful requests, or raise an alarm when errors exceed a threshold.

CloudWatch Metrics for Amazon Lex

To get metrics for your Amazon Lex operations, you must specify the following information:

- The metric dimension. A *dimension* is a set of name-value pairs that you use to identify a metric. Amazon Lex has three dimensions:
 - BotAlias, BotName, Operation
 - BotAlias, BotName, InputMode, Operation
 - BotName, BotVersion, InputMode, Operation
- The metric name, such as MissedUtteranceCount or RuntimeRequestCount.

You can get metrics for Amazon Lex with the AWS Management Console, the AWS CLI, or the CloudWatch API. You can use the CloudWatch API through one of the Amazon AWS Software Development Kits (SDKs) or the CloudWatch API tools. The Amazon Lex console displays graphs based on the raw data from the CloudWatch API.

You must have the appropriate CloudWatch permissions to monitor Amazon Lex with CloudWatch. For more information, see [Authentication and Access Control for Amazon CloudWatch](#) in the *Amazon CloudWatch User Guide*.

Viewing Amazon Lex Metrics

View Amazon Lex metrics using the Amazon Lex console or the CloudWatch console.

To view metrics (Amazon Lex console)

1. Sign in to the AWS Management Console and open the Amazon Lex console at <https://console.aws.amazon.com/lex/>.
2. From the list of bots, choose the one whose metrics you want to see.
3. Choose **Monitoring**. Metrics are displayed in graphs.

To view metrics (CloudWatch console)

1. Sign in to the AWS Management Console and open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Metrics**, choose **All Metrics**, and then choose **AWS/Lex**.
3. Choose the dimension, choose a metric name, then choose **Add to graph**.
4. Choose a value for the date range. The metric count for the selected date range is displayed in the graph.

Creating an Alarm

A CloudWatch alarm watches a single metric over a specified time period, and performs one or more actions: sending a notification to an Amazon Simple Notification Service (Amazon SNS) topic or Auto Scaling policy. The action or actions are based on the value of the metric relative to a given threshold over a number of time periods that you specify. CloudWatch can also send you an Amazon SNS message when the alarm changes state.

CloudWatch alarms invoke actions only when the state changes and has persisted for the period that you specify.

To set an alarm

1. Sign in to the AWS Management Console and open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Alarms**, and then choose **Create Alarm**.
3. Choose **AWS/Lex Metrics**, and then choose a metric.
4. For **Time Range**, choose a time range to monitor, and then choose **Next**.
5. Enter a **Name** and **Description**.
6. For **Whenever**, choose **>=**, and type a maximum value.

7. If you want CloudWatch to send an email when the alarm state is reached, in the **Actions** section, for **Whenever this alarm**, choose **State is ALARM**. For **Send notification to**, choose a mailing list or choose **New list** and create a new mailing list.
8. Preview the alarm in the **Alarm Preview** section. If you are satisfied with the alarm, choose **Create Alarm**.

CloudWatch Metrics for Amazon Lex Runtime

The following table describes the Amazon Lex runtime metrics.

Metric	Description
KendraIndexAccessError	The number of times that Amazon Lex could not access your Amazon Kendra index. Valid dimension for the PostContent operation with the Text or Speech InputMode : • BotName, BotAlias, Operation, InputMode Valid dimension for the PostText operation: • BotName, BotAlias, Operation Unit: Count
KendraLatency	The amount of time that it takes Amazon Kendra to respond to a request from the AMAZON.KendraSearchIntent . Valid dimensions for the PostContent operation with the Text or Speech InputMode : • BotName, BotVersion, Operation, InputMode • BotName, BotAlias, Operation, InputMode Valid dimensions for the PostText operation:

Metric	Description
	• BotName, BotVersion, Operation • BotName, BotAlias, Operation Unit: Milliseconds
KendraSuccess	The number of successful requests from the AMAZON.KendraSearchIntent to your Amazon Kendra index. Valid dimensions for the PostContent operation with the Text or Speech InputMode : • BotName, BotVersion, Operation, InputMode • BotName, BotAlias, Operation, InputMode Valid dimensions for the PostText operation: • BotName, BotVersion, Operation • BotName, BotAlias, Operation Unit: Count
KendraSystemErrors	The number of times that Amazon Lex couldn't query the Amazon Kendra index. Valid dimension for the PostContent operation with the Text or Speech InputMode : • BotName, BotAlias, Operation, InputMode Valid dimension for the PostText operation: • BotName, BotAlias, Operation Unit: Count

Metric	Description
KendraThrottledEvents	The number of times Amazon Kendra throttled requests from the <code>AMAZON.KendraSearchIntent</code> . Valid dimension for the <code>PostContent</code> operation with the <code>Text</code> or <code>Speech</code> <code>InputMode</code> : • <code>BotName</code>, <code>BotAlias</code>, <code>Operation</code>, <code>InputMode</code> Valid dimension for the <code>PostText</code> operation: • <code>BotName</code>, <code>BotAlias</code>, <code>Operation</code> Unit: Count
MissedUtteranceCount	The number of utterances that were not recognized in the specified period. Valid dimensions for the <code>PostContent</code> operation with the <code>Text</code> or <code>Speech</code> <code>InputMode</code> : • <code>BotName</code>, <code>BotVersion</code>, <code>Operation</code>, <code>InputMode</code> • <code>BotName</code>, <code>BotAlias</code>, <code>Operation</code>, <code>InputMode</code> Valid dimensions for the <code>PostText</code> operation: • <code>BotName</code>, <code>BotVersion</code>, <code>Operation</code> • <code>BotName</code>, <code>BotAlias</code>, <code>Operation</code>

Metric	Description
RuntimeConcurrency	The number of concurrent connections in the specified time period. RuntimeConcurrency is reported as a <code>StatisticSet</code> . Valid dimensions for the <code>PostContent</code> operation with the <code>Text</code> or <code>Speech</code> <code>InputMode</code> : • Operation, BotName, BotVersion, InputMode • Operation, BotName, BotAlias, InputMode Valid dimensions for other operations: • Operation, BotName, BotVersion • Operation, BotName, BotAlias Unit: Count
RuntimeInvalidLambdaResponses	The number of invalid AWS Lambda (Lambda) responses in the specified period. Valid dimension for the <code>PostContent</code> operation with the <code>Text</code> or <code>Speech</code> <code>InputMode</code> : • BotName, BotAlias, Operation, InputMode Valid dimension for the <code>PostText</code> operation: • BotName, BotAlias, Operation

Metric	Description
RuntimeLambdaErrors	The number of Lambda runtime errors in the specified period. Valid dimension for the PostContent operation with the Text or Speech InputMode : • BotName, BotAlias, Operation, InputMode Valid dimension for the PostText operation: • BotName, BotAlias, Operation
RuntimePollyErrors	The number of invalid Amazon Polly responses in the specified period. Valid dimension for the PostContent operation with the Text or Speech InputMode : • BotName, BotAlias, Operation, InputMode Valid dimension for the PostText operation: • BotName, BotAlias, Operation

Metric	Description
RuntimeRequestCount	The number of runtime requests in the specified period. Valid dimensions for the PostContent operation with the Text or Speech InputMode : • BotName, BotVersion, Operation, InputMode • BotName, BotAlias, Operation, InputMode Valid dimensions for the PostText operation: • BotName, BotVersion, Operation • BotName, BotAlias, Operation Unit: Count
 Important This metric is RuntimeSuccessfulRequestLatency and not RuntimeSuccessfulRequestLatency .	The latency for successful requests between the time that the request was made and the response was passed back. Valid dimensions for the PostContent operation with the Text or Speech InputMode : • BotName, BotVersion, Operation, InputMode • BotName, BotAlias, Operation, InputMode Valid dimensions for the PostText operation: • BotName, BotVersion, Operation • BotName, BotAlias, Operation Unit: Milliseconds

Metric	Description
RuntimeSystemErrors	The number of system errors in the specified period. The response code range for a system error is 500 to 599. Valid dimension for the PostContent operation with the Text or Speech InputMode : • BotName, BotAlias, Operation, InputMode Valid dimension for the PostText operation: • BotName, BotAlias, Operation Unit: Count
RuntimeThrottledEvents	The number of throttled requests. Amazon Lex throttles a request when it receives more requests than the limit of transactions per second set for your account. If the limit set for your account is frequently exceeded, you can request a limit increase. To request an increase, see AWS Service Limits. Valid dimension for the PostContent operation with the Text or Speech InputMode : • BotName, BotAlias, Operation, InputMode Valid dimension for the PostText operation: • BotName, BotAlias, Operation Unit: Count

Metric	Description
RuntimeUserErrors	The number of user errors in the specified period. The response code range for a user error is 400 to 499. Valid dimension for the PostContent operation with Text or Speech InputMode : • BotName, BotAlias, Operation, InputMode Valid dimension for the PostText operation: • BotName, BotAlias, Operation Unit: Count

Amazon Lex runtime metrics use the AWS/Lex namespace, and provide metrics in the following dimensions. You can group metrics by dimensions in the CloudWatch console:

Dimension	Description
BotName, BotAlias, Operation, InputMode	Groups metrics by the bot's alias, the bot's name, the operation (PostContent), and by whether the input was text or speech.
BotName, BotVersion, Operation, InputMode	Groups metrics by the bot's name, the version of the bot, the operation (PostContent), and by whether the input was text or speech.
BotName, BotVersion, Operation	Groups metrics by the bot's name, the bots version, and by the operation, PostText.
BotName, BotAlias, Operation	Groups metrics by the bot's name, the bot's alias, and by the operation, PostText.

CloudWatch Metrics for Amazon Lex Channel Associations

A channel association is the association between Amazon Lex and a messaging channel, such as Facebook. The following table describes the Amazon Lex channel association metrics.

Metric	Description
BotChannelAuthErrors	The number of authentication errors returned by the messaging channel in the specified time period. An authentication error indicates that the secret token provided during channel creation is invalid or has expired.
BotChannelConfigurationErrors	The number of configuration errors in the specified period. A configuration error indicates that one or more configuration entries for the channel are invalid.
BotChannelInboundThrottledEvents	The number of times that messages that were sent by the messaging channel were throttled by Amazon Lex in the specified period.
BotChannelOutboundThrottledEvents	The number of times that outbound events from Amazon Lex to the messaging channel were throttled in the specified time period.
BotChannelRequestCount	The number of requests made on a channel in the specified time period.
BotChannelResponseCardErrors	The number of times that Amazon Lex could not post response cards in the specified period.
BotChannelSystemErrors	The number of internal errors that occurred in Amazon Lex for a channel in the specified period.

Amazon Lex channel association metrics use the `AWS/Lex` namespace, and provide metrics for the following dimension. You can group metrics by dimensions in the CloudWatch console:

Dimension	Description
BotAlias, BotChannelName, BotName, Source	Group metrics by the bot's alias, the channel name, the bot's name, and the source of traffic.

CloudWatch Metrics for Conversation Logs

Amazon Lex uses the following metrics for conversation logging:

Metric	Description
ConversationLogsAudioDeliverySuccess	The number of audio logs successfully delivered to the S3 bucket in the specified time period. Units: Count
ConversationLogsAudioDeliveryFailure	The number of audio logs that failed to be delivered to the S3 bucket in the specified time period. A delivery failure indicates an error with the resources configured for conversation logs. Errors can include insufficient IAM permissions, an inaccessible AWS KMS key, or an inaccessible S3 bucket. Units: Count
ConversationLogsTextDeliverySuccess	The number of text logs successfully delivered to CloudWatch Logs in the specified time period. Units: Count
ConversationLogsTextDeliveryFailure	The number of text logs that failed to be delivered to CloudWatch Logs in the specified time period. A delivery failure indicates an error with the resources configured for

Metric	Description
	conversation logs. Errors can include insufficient IAM permissions, an inaccessible AWS KMS key, or an inaccessible CloudWatch Logs log group. Units: Count

Amazon Lex conversation log metrics use the AWS/Lex namespace, and provide metrics for the following dimensions. You can group metrics by dimension in the CloudWatch console.

Dimension	Description
BotAlias	Group metrics by the bot's alias.
BotName	Group metrics by the bot's name.
BotVersion	Group metrics by the bot's version.

Monitoring Amazon Lex API Calls with AWS CloudTrail Logs

Amazon Lex is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Amazon Lex. CloudTrail captures a subset of API calls for Amazon Lex as events, including calls from the Amazon Lex console and from code calls to the Amazon Lex APIs. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Lex. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon Lex, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, including how to configure and enable it, see the [AWS CloudTrail User Guide](#).

Amazon Lex Information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When supported event activity occurs in Amazon Lex, that activity is recorded in a CloudTrail event along with other AWS

service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your AWS account, including events for Amazon Lex, create a trail. A trail enables CloudTrail to deliver log files to an Amazon Simple Storage Service (Amazon S3) bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)

Amazon Lex supports logging the following operations as events in CloudTrail log files:

- [CreateBotVersion](#)
- [CreateIntentVersion](#)
- [CreateSlotTypeVersion](#)
- [DeleteBot](#)
- [DeleteBotAlias](#)
- [DeleteBotChannelAssociation](#)
- [DeleteBotVersion](#)
- [DeleteIntent](#)
- [DeleteIntentVersion](#)
- [DeleteSlotType](#)
- [DeleteSlotTypeVersion](#)
- [DeleteUtterances](#)
- [GetBot](#)
- [GetBotAlias](#)
- [GetBotAliases](#)
- [GetBotChannelAssociation](#)

- [GetBotChannelAssociations](#)
- [GetBots](#)
- [GetBotVersions](#)
- [GetBuiltinIntent](#)
- [GetBuiltinIntents](#)
- [GetBuiltinSlotTypes](#)
- [GetSlotTypeVersions](#)
- [GetUtterancesView](#)
- [PutBot](#)
- [PutBotAlias](#)
- [PutIntent](#)
- [PutSlotType](#)

Every event or log entry contains information about who generated the request. This information helps you determine the following:

- Whether the request was made with root or user credentials
- Whether the request was made with temporary security credentials for a role or federated user
- Whether the request was made by another AWS service

For more information, see the [CloudTrail userIdentity Element](#).

For information about the Amazon Lex actions that are logged in CloudTrail logs, see [Amazon Lex Model Building Service](#). For example, calls to the [PutBot](#), [GetBot](#), and [DeleteBot](#) operations generate entries in the CloudTrail log. The actions documented in [Amazon Lex Runtime Service](#), [PostContent](#) and [PostText](#), are not logged.

Example: Amazon Lex Log File Entries

A trail is a configuration that enables delivery of events as log files to an S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files are not an ordered stack trace of the public API calls, so they do not appear in any specific order.

The following example CloudTrail log entry shows the result of a call to the PutBot operation.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "AssumedRole | FederatedUser | IAMUser | Root | SAMLUser | WebIdentityUser",
    "principalId": "principal ID",
    "arn": "ARN",
    "accountId": "account ID",
    "accessKeyId": "access key ID",
    "userName": "user name"
  },
  "eventTime": "timestamp",
  "eventSource": "lex.amazonaws.com",
  "eventName": "PutBot",
  "awsRegion": "region",
  "sourceIPAddress": "source IP address",
  "userAgent": "user agent",
  "requestParameters": {
    "name": "CloudTrailBot",
    "intents": [
      {
        "intentVersion": "11",
        "intentName": "TestCloudTrail"
      }
    ],
    "voiceId": "Salli",
    "childDirected": false,
    "locale": "en-US",
    "idleSessionTTLInSeconds": 500,
    "processBehavior": "BUILD",
    "description": "CloudTrail test bot",
    "clarificationPrompt": {
      "messages": [
        {
          "contentType": "PlainText",
          "content": "I didn't understand you. What would you
like to do?"
        }
      ],
      "maxAttempts": 2
    },
    "abortStatement": {
```

```

        "messages": [
            {
                "contentType": "PlainText",
                "content": "Sorry. I'm not able to assist at this
time."
            }
        ]
    },
    "responseElements": {
        "voiceId": "Salli",
        "locale": "en-US",
        "childDirected": false,
        "abortStatement": {
            "messages": [
                {
                    "contentType": "PlainText",
                    "content": "Sorry. I'm not able to assist at this
time."
                }
            ]
        },
        "status": "BUILDING",
        "createdAt": "timestamp",
        "lastUpdatedAt": "timestamp",
        "idleSessionTTLInSeconds": 500,
        "intents": [
            {
                "intentVersion": "11",
                "intentName": "TestCloudTrail"
            }
        ],
        "clarificationPrompt": {
            "messages": [
                {
                    "contentType": "PlainText",
                    "content": "I didn't understand you. What would you
like to do?"
                }
            ]
        },
        "maxAttempts": 2
    },
    "version": "$LATEST",
    "description": "CloudTrail test bot",

```



```
        "checksum": "checksum",
        "name": "CloudTrailBot"
    },
    "requestID": "request ID",
    "eventID": "event ID",
    "eventType": "AwsApiCall",
    "recipientAccountId": "account ID"
}
}
```

Compliance Validation for Amazon Lex

Third-party auditors assess the security and compliance of Amazon Lex as part of multiple AWS compliance programs. Amazon Lex is a HIPAA eligible service. It is PCI, SOC, and ISO compliant. You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using Amazon Lex is determined by the sensitivity of your data, your organization's compliance objectives, and applicable laws and regulations. If your use of Amazon Lex is subject to compliance with standards such as PCI, AWS provides the following resources to help:

- [Security and Compliance Quick Start Guides](#) – Deployment guides that discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS
- [Architecting for HIPAA Security and Compliance Whitepaper](#) – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.
- [AWS Compliance Resources](#) – A collection of workbooks and guides that might apply to your industry and location
- [AWS Config](#) – A service that assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations
- [AWS Security Hub](#) – A comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices

For a list of AWS services in scope for specific compliance programs, see [AWS Services in Scope by Compliance Program](#). For general information, see [AWS Compliance Programs](#).

Resilience in Amazon Lex

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

In addition to the AWS global infrastructure, Amazon Lex offers several features to help support your data resiliency and backup needs.

Infrastructure Security in Amazon Lex

As a managed service, Amazon Lex is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

You use published AWS API calls to access Amazon Lex through the network. Clients must support TLS (Transport Layer Security) 1.0. We recommend TLS 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS), such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems, such as Java 7 and later, support these modes. Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

You can call these API operations from any network location, but Amazon Lex supports resource-level access policies, which can include restrictions based on the source IP address. You can also use Amazon Lex policies to control access from specific Amazon Virtual Private Cloud (Amazon VPC) endpoints or specific VPCs. Effectively, this isolates network access to a given Amazon Lex resource from only the specific VPC within the AWS network.

Guidelines and Quotas in Amazon Lex

The following sections provide guidelines and quotas when using Amazon Lex.

Topics

- [Supported Regions](#)
- [General Guidelines](#)
- [Quotas](#)

Supported Regions

For a list of AWS Regions where Amazon Lex is available, see [AWS Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

General Guidelines

This section describes general guidelines when using Amazon Lex.

- Signing requests – All Amazon Lex model-building and runtime API operations in the [API Reference](#) use signature V4 for authenticating requests. For more information about authenticating requests, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*.

For [PostContent](#), Amazon Lex uses the unsigned payload option described in [Signature Calculations for the Authorization Header: Transferring Payload in a Single Chunk \(AWS Signature Version 4\)](#) in the *Amazon Simple Storage Service (S3) API Reference*.

When you use the unsigned payload option, don't include the hash of the payload in the canonical request. Instead, you use the literal string "UNSIGNED-PAYLOAD" as the hash of the payload. Also include a header with the name `x-amz-content-sha256` and the value `UNSIGNED-PAYLOAD` in the `PostContent` request.

- Note the following about how Amazon Lex captures slot values from user utterances:

Amazon Lex uses the enumeration values you provide in a slot type definition to train its machine learning models. Suppose you define an intent called `GetPredictionIntent` with the following sample utterance:

```
"Tell me the prediction for {Sign}"
```

Where `{Sign}` is a slot of custom type `ZodiacSign`. It has 12 enumeration values, Aries through Pisces. From the user utterance "Tell me the prediction for ..." Amazon Lex understands what follows is a zodiac sign.

When the `valueSelectionStrategy` field is set to `ORIGINAL_VALUE` using the [PutSlotType](#) operation, or if **Expand values** is selected in the console, if the user says "Tell me the prediction for earth", Amazon Lex infers that "earth" is a `ZodiacSign` and passes it to your client application or Lambda functions. You must check that slot values have valid values before using them in your fulfillment activity.

If you set the `valueSelectionStrategy` field to `TOP_RESOLUTION` using the [PutSlotType](#) operation, or if **Restrict to slot values and synonyms** is selected in the console, the values that are returned are limited to the values that you defined for the slot type. For example, if the user says "Tell me the prediction for earth" the value would not be recognized because it is not one of the values defined for the slot type. When you define synonyms for slot values, they are recognized the same as a slot value, however, the slot value is returned instead of the synonym.

When Amazon Lex calls a Lambda function or returns the result of a speech interaction with your client application, the case of the slot values is not guaranteed. For example, if you are eliciting values for the [AMAZON.Movie](#) built-in slot type, and a user says or types "Gone with the wind," Amazon Lex may return "Gone with the Wind," "gone with the wind," or "Gone With The Wind." In text interactions, the case of the slot values matches the text entered or the slot value, depending on the value of the `valueResolutionStrategy` field.

- When defining slot values that contain acronyms, use the following patterns:
 - Capital letters separated by periods (D.V.D.)
 - Capital letters separated by spaces (D V D)
- Amazon Lex does not support the AMAZON.LITERAL built-in slot type that the Alexa Skills Kit supports. However, Amazon Lex supports creating custom slot types that you can use to implement this functionality. As mentioned in the previous bullet, you can capture values outside the custom slot type definition. Add more and diverse enumeration values to boost the automatic speech recognition (ASR) and natural language understanding (NLU) accuracy.
- The [AMAZON.DATE](#) and [AMAZON.TIME](#) built-in slot types capture both absolute and relative dates and times. Relative dates and times are resolved in the region where Amazon Lex is processing the request.

For the AMAZON.TIME built-in slot type, if the user doesn't specify that a time is before or after noon, the time is ambiguous and Amazon Lex will prompt the user again. We recommend prompts that elicit an absolute time. For example, use a prompt such as "When do you want your pizza delivered? You can say 6 PM or 6 in the evening."

- Providing confusable training data in your bot reduces Amazon Lex's ability to understand user input. Consider these examples:

Suppose you have two intents (OrderPizza and OrderDrink) in your bot and both are configured with an "I want to order" utterance. This utterance does not map to a specific intent that Amazon Lex can learn from while building the language model for the bot at build time. As a result, when a user inputs this utterance at runtime, Amazon Lex can't pick an intent with a high degree of confidence.

Consider another example where you define a custom intent for getting a confirmation from the user (for example, `MyCustomConfirmationIntent`) and configure the intent with the utterances "Yes" and "No." Note that Amazon Lex also has a language model for understanding user confirmations. This can create conflicting situation. When the user responds with a "Yes," does this mean that this is a confirmation for the ongoing intent or that the user is requesting the custom intent that you created?

In general, the sample utterances you provide should map to a specific intent and, optionally, to specific slot values.

- The runtime API operations [PostContent](#) and [PostText](#) take a user ID as the required parameter. Developers can set this to any value that meets the constraints described in the API. We recommend you don't use this parameter to send any confidential information such as user logins, emails, or social security numbers. This ID is primarily used to uniquely identify conversation with a bot (there can be multiple users ordering pizza).
- If your client application uses Amazon Cognito for authentication, you might use the Amazon Cognito user ID as Amazon Lex user ID. Note that any Lambda function configured for your bot must have its own authentication mechanism to identify the user on whose behalf Amazon Lex is invoking the Lambda function.
- We encourage you to define an intent that captures a user's intention to discontinue the conversation. For example, you can define an intent (`NothingIntent`) with sample utterances ("I don't want anything", "exit", "bye bye"), no slots, and no Lambda function configured as a code hook. This lets users gracefully close a conversation.

Quotas

This section describes current quotas in Amazon Lex. These quotas are grouped by categories.

Service quotas can be adjusted or increased. Contact AWS customer support to increase a quota. It can take a few days to increase a service quota. If you're increasing your quota as part of a larger project, be sure to add this time to your plan.

Topics

- [Runtime Service Quotas](#)
- [Model Building Quotas](#)

Runtime Service Quotas

In addition to the quotas described in the API reference, note the following:

API Quotas

- Speech input to the [PostContent](#) operation can be up to 15 seconds long.
- In both the runtime API operations [PostContent](#) and [PostText](#), the input text size can be up to 1024 Unicode characters.
- The maximum size of PostContent headers is 16 KB. The maximum size of request and session headers combined is 12 KB.
- When using the PostContent or PostText operations in text mode, the maximum number of concurrent conversations with a bot is 2 for the \$LATEST alias and 50 for all other aliases. The quota applies separately for each API.
- When using the PostContent operation in voice mode, the maximum number of concurrent text-mode conversations with a bot is 2 for the \$LATEST alias and 125 for all other aliases. The quota applies separately for each API.
- The maximum number of concurrent session management calls ([PutSession](#), [GetSession](#), and [DeleteSession](#)) is 2 for the \$LATEST alias of a bot and 50 for all other aliases.

- The maximum input size to a Lambda function is 12 KB. The maximum output size is 25 KB, of which 12 KB can be session attributes.

Using the \$LATEST version

- The \$LATEST version of your bot should only be used for manual testing. Amazon Lex limits the number of runtime requests that you can make to the \$LATEST version of the bot.
- When you update the \$LATEST version of the bot, Amazon Lex terminates any in-progress conversations for any client application using the \$LATEST version of the bot. Generally, you should not use the \$LATEST version of a bot in production because \$LATEST version can be updated. You should publish a version and use it instead.
- When you update an alias, Amazon Lex takes a few minutes to pick up the change. When you modify the \$LATEST version of the bot, the change is picked up immediately.

Session Timeout

- The session timeout set when the bot was created determines how long the bot retains conversation context, such as current user intent and slot data.
- After a user starts the conversation with your bot and until the session expires, Amazon Lex uses the same bot version, even if you update the bot alias to point to another version.

Model Building Quotas

Model building refers to creating and managing bots. This includes creating and managing bots, intents, slot types, slots, and bot channel associations.

Topics

- [Bot Quotas](#)
- [Intent Quotas](#)
- [Slot Type Quotas](#)

Bot Quotas

- You configure prompts and statements throughout the model building API. Each of these prompts or statements can have up to five messages and each message can contain from 1 to 1000 UTF-8 characters.
- When using message groups you can define up to five message groups for each message. Each message group can contain a maximum of five messages, and you are limited to 15 messages in all message groups.
- You can define sample utterances for intents and slots. You can use a maximum of 200,000 characters for all utterances.
- Each slot type can define a maximum of 10,000 values and synonyms. Each bot can contain a maximum of 50,000 slot type values and synonyms.
- Bot, alias, and bot channel association names are case insensitive at the time of creation. If you create PizzaBot and then try to create pizzaBot, you will get an error. However, when accessing a resource, the resource names are case sensitive, you must specify PizzaBot and not pizzaBot. These names must be between 2 and 50 ASCII characters.

- The maximum number of versions you can publish for all resource types is 100. Note that there is no versioning for aliases.
- Within a bot, intent names and slot names must be unique, you can't have an intent and a slot by the same name.
- You can create a bot that is configured to support multiple intents. If two intents have a slot by the same name, then the corresponding slot type must be the same.

For example, suppose you create a bot to support two intents (`OrderPizza` and `OrderDrink`). If both these intents have the `size` slot, then the slot type must be the same in both places.

In addition, the sample utterances you provide for a slot in one of the intents applies to a slot with the same name in other intents.

- You can associate a maximum of 250 intents with a bot.
- When you create a bot, you specify a session timeout. The session timeout can be between one minute and one day. The default is five minutes.
- You can create up to five aliases for a bot.
- You can create up to 250 bots per AWS account.
- You cannot create multiple intents that extend from the same built-in intent.

Intent Quotas

- Intent and slot names are case insensitive at the time of creation. That is, if you create `OrderPizza` intent and then again try to create another `orderPizza` intent, you will get an error. However, when accessing these resources, the resource names are case sensitive, specify `OrderPizza` and not `orderPizza`. These names must be between 1 and 100 ASCII characters.
- An intent can have up to 1,500 sample utterances. A minimum of one sample utterance is required. Each sample utterance can be up to 200 UTF-8 characters long. You can use up to 200,000 characters for all intent and slot utterances in a bot. A sample utterance for an intent:
 - Can refer to zero or more slot names.
 - Can refer to a slot name only once.

For example:

```
I want a pizza  
I want a {pizzaSize} pizza  
I want a {pizzaSize} {pizzaTopping} pizza
```

- Although each intent supports up to 1,500 utterances, if you use fewer utterances Amazon Lex may have a better ability to recognize inputs outside your provided set.
- You can create up to five message groups for each message in an intent. There can be a total of 15 messages in all message groups for a message.
- The console can only create message groups for the `conclusionStatement` and `followUpPrompt` messages. You can create message groups for any other message using the Amazon Lex API.
- Each slot can have up to 10 sample utterances. Each sample utterance must refer to the slot name exactly once. For example:

```
{pizzaSize} please
```

- Each bot can have a maximum of 200,000 characters for intent and slot utterances combined.
- You cannot provide utterances for intents that extend from built-in intents. For all other intents you must provide at least one sample utterance. Intents contain slots, but the slot level sample utterances are optional.
- Built-in intents
 - Currently, Amazon Lex does not support slot elicitation for built-in intents. You cannot create Lambda functions to return the `ElicitSlot` directive in the response with an intent that is derived from built-in intents. For more information, see [Response Format](#).
 - The service does not support adding sample utterances to built-in intents. Similarly, you cannot add or remove slots to built-in intents.
- You can create up to 1,000 intents per AWS account. You can create up to 100 slots in an intent.

Slot Type Quotas

- Slot type names are case insensitive at the time of creation. If you create the `PizzaSize` slot type and then again try to create the `pizzaSize` slot type, you will get an error. However, when accessing these resources, the resource names are case sensitive (you must specify `PizzaSize` and not `pizzaSize`). Names must be between 1 and 100 ASCII characters.
- A custom slot type you create can have a maximum of 10,000 enumeration values and synonyms. Each value can be up to 140 UTF-8 characters long. The enumeration values and synonyms cannot contain duplicates.

- For a slot type value, where appropriate, specify both upper and lower case. For example, for a slot type called Procedure, if value is MRI, specify both "MRI" and "mri" as values.
- Built-in slot types – Currently, Amazon Lex doesn't support adding enumeration values or synonyms for the built-in slot types.

API Reference

This section provides documentation for the Amazon Lex API operations. For a list of AWS Regions where Amazon Lex is available, see [AWS Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

Topics

- [Actions](#)
- [Data Types](#)

Actions

The following actions are supported by Amazon Lex Model Building Service:

- [CreateBotVersion](#)
- [CreateIntentVersion](#)
- [CreateSlotTypeVersion](#)
- [DeleteBot](#)
- [DeleteBotAlias](#)
- [DeleteBotChannelAssociation](#)
- [DeleteBotVersion](#)
- [DeleteIntent](#)
- [DeleteIntentVersion](#)
- [DeleteSlotType](#)
- [DeleteSlotTypeVersion](#)
- [DeleteUtterances](#)
- [GetBot](#)
- [GetBotAlias](#)
- [GetBotAliases](#)
- [GetBotChannelAssociation](#)
- [GetBotChannelAssociations](#)
- [GetBots](#)

- [GetBotVersions](#)
- [GetBuiltinIntent](#)
- [GetBuiltinIntents](#)
- [GetBuiltinSlotTypes](#)
- [GetExport](#)
- [GetImport](#)
- [GetIntent](#)
- [GetIntents](#)
- [GetIntentVersions](#)
- [GetMigration](#)
- [GetMigrations](#)
- [GetSlotType](#)
- [GetSlotTypes](#)
- [GetSlotTypeVersions](#)
- [GetUtterancesView](#)
- [ListTagsForResource](#)
- [PutBot](#)
- [PutBotAlias](#)
- [PutIntent](#)
- [PutSlotType](#)
- [StartImport](#)
- [StartMigration](#)
- [TagResource](#)
- [UntagResource](#)

The following actions are supported by Amazon Lex Runtime Service:

- [DeleteSession](#)
- [GetSession](#)
- [PostContent](#)
- [PostText](#)

- [PutSession](#)

Amazon Lex Model Building Service

The following actions are supported by Amazon Lex Model Building Service:

- [CreateBotVersion](#)
- [CreateIntentVersion](#)
- [CreateSlotTypeVersion](#)
- [DeleteBot](#)
- [DeleteBotAlias](#)
- [DeleteBotChannelAssociation](#)
- [DeleteBotVersion](#)
- [DeleteIntent](#)
- [DeleteIntentVersion](#)
- [DeleteSlotType](#)
- [DeleteSlotTypeVersion](#)
- [DeleteUtterances](#)
- [GetBot](#)
- [GetBotAlias](#)
- [GetBotAliases](#)
- [GetBotChannelAssociation](#)
- [GetBotChannelAssociations](#)
- [GetBots](#)
- [GetBotVersions](#)
- [GetBuiltinIntent](#)
- [GetBuiltinIntents](#)
- [GetBuiltinSlotTypes](#)
- [GetExport](#)
- [GetImport](#)
- [GetIntent](#)
- [GetIntents](#)

- [GetIntentVersions](#)
- [GetMigration](#)
- [GetMigrations](#)
- [GetSlotType](#)
- [GetSlotTypes](#)
- [GetSlotTypeVersions](#)
- [GetUtterancesView](#)
- [ListTagsForResource](#)
- [PutBot](#)
- [PutBotAlias](#)
- [PutIntent](#)
- [PutSlotType](#)
- [StartImport](#)
- [StartMigration](#)
- [TagResource](#)
- [UntagResource](#)

CreateBotVersion

Service: Amazon Lex Model Building Service

Creates a new version of the bot based on the `$LATEST` version. If the `$LATEST` version of this resource hasn't changed since you created the last version, Amazon Lex doesn't create a new version. It returns the last created version.

Note

You can update only the `$LATEST` version of the bot. You can't update the numbered versions that you create with the `CreateBotVersion` operation.

When you create the first version of a bot, Amazon Lex sets the version to 1. Subsequent versions increment by 1. For more information, see [Versioning](#).

This operation requires permission for the `lex:CreateBotVersion` action.

Request Syntax

```
POST /bots/name/versions HTTP/1.1
Content-type: application/json

{
  "checksum": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the bot that you want to create a new version of. The name is case sensitive.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request accepts the following data in JSON format.

checksum

Identifies a specific revision of the \$LATEST version of the bot. If you specify a checksum and the \$LATEST version of the bot has a different checksum, a `PreconditionFailedException` exception is returned and Amazon Lex doesn't publish a new version. If you don't specify a checksum, Amazon Lex publishes the \$LATEST version.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 201
Content-type: application/json

{
  "abortStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "checksum": "string",
  "childDirected": boolean,
  "clarificationPrompt": {
    "maxAttempts": number,
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ]
  },
}
```

```

    "responseCard": "string"
  },
  "createdDate": number,
  "description": "string",
  "detectSentiment": boolean,
  "enableModelImprovements": boolean,
  "failureReason": "string",
  "idleSessionTTLInSeconds": number,
  "intents": [
    {
      "intentName": "string",
      "intentVersion": "string"
    }
  ],
  "lastUpdatedDate": number,
  "locale": "string",
  "name": "string",
  "status": "string",
  "version": "string",
  "voiceId": "string"
}

```

Response Elements

If the action is successful, the service sends back an HTTP 201 response.

The following data is returned in JSON format by the service.

abortStatement

The message that Amazon Lex uses to cancel a conversation. For more information, see [PutBot](#).

Type: [Statement](#) object

checksum

Checksum identifying the version of the bot that was created.

Type: String

childDirected

For each Amazon Lex bot created with the Amazon Lex Model Building Service, you must specify whether your use of Amazon Lex is related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to

the Children's Online Privacy Protection Act (COPPA) by specifying `true` or `false` in the `childDirected` field. By specifying `true` in the `childDirected` field, you confirm that your use of Amazon Lex **is** related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA. By specifying `false` in the `childDirected` field, you confirm that your use of Amazon Lex **is not** related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA. You may not specify a default value for the `childDirected` field that does not accurately reflect whether your use of Amazon Lex is related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA.

If your use of Amazon Lex relates to a website, program, or other application that is directed in whole or in part, to children under age 13, you must obtain any required verifiable parental consent under COPPA. For information regarding the use of Amazon Lex in connection with websites, programs, or other applications that are directed or targeted, in whole or in part, to children under age 13, see the [Amazon Lex FAQ](#).

Type: Boolean

[**clarificationPrompt**](#)

The message that Amazon Lex uses when it doesn't understand the user's request. For more information, see [PutBot](#).

Type: [Prompt](#) object

[**createdDate**](#)

The date when the bot version was created.

Type: Timestamp

[**description**](#)

A description of the bot.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

[**detectSentiment**](#)

Indicates whether utterances entered by the user should be sent to Amazon Comprehend for sentiment analysis.

Type: Boolean

[enableModelImprovements](#)

Indicates whether the bot uses accuracy improvements. `true` indicates that the bot is using the improvements, otherwise, `false`.

Type: Boolean

[failureReason](#)

If status is `FAILED`, Amazon Lex provides the reason that it failed to build the bot.

Type: String

[idleSessionTTLInSeconds](#)

The maximum time in seconds that Amazon Lex retains the data gathered in a conversation. For more information, see [PutBot](#).

Type: Integer

Valid Range: Minimum value of 60. Maximum value of 86400.

[intents](#)

An array of Intent objects. For more information, see [PutBot](#).

Type: Array of [Intent](#) objects

[lastUpdatedDate](#)

The date when the `$LATEST` version of this bot was updated.

Type: Timestamp

[locale](#)

Specifies the target locale for the bot.

Type: String

Valid Values: `de-DE` | `en-AU` | `en-GB` | `en-IN` | `en-US` | `es-419` | `es-ES` | `es-US` | `fr-FR` | `fr-CA` | `it-IT` | `ja-JP` | `ko-KR`

[name](#)

The name of the bot.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

status

When you send a request to create or update a bot, Amazon Lex sets the status response element to BUILDING. After Amazon Lex builds the bot, it sets status to READY. If Amazon Lex can't build the bot, it sets status to FAILED. Amazon Lex returns the reason for the failure in the failureReason response element.

Type: String

Valid Values: BUILDING | READY | READY_BASIC_TESTING | FAILED | NOT_BUILT

version

The version of the bot.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

voiceId

The Amazon Polly voice ID that Amazon Lex uses for voice interactions with the user.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

PreconditionFailedException

The checksum of the resource that you are trying to change does not match the checksum in the request. Check the resource's checksum and try again.

HTTP Status Code: 412

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

CreateIntentVersion

Service: Amazon Lex Model Building Service

Creates a new version of an intent based on the \$LATEST version of the intent. If the \$LATEST version of this intent hasn't changed since you last updated it, Amazon Lex doesn't create a new version. It returns the last version you created.

Note

You can update only the \$LATEST version of the intent. You can't update the numbered versions that you create with the CreateIntentVersion operation.

When you create a version of an intent, Amazon Lex sets the version to 1. Subsequent versions increment by 1. For more information, see [Versioning](#).

This operation requires permissions to perform the `lex:CreateIntentVersion` action.

Request Syntax

```
POST /intents/name/versions HTTP/1.1
Content-type: application/json

{
  "checksum": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the intent that you want to create a new version of. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

Request Body

The request accepts the following data in JSON format.

checksum

Checksum of the \$LATEST version of the intent that should be used to create the new version. If you specify a checksum and the \$LATEST version of the intent has a different checksum, Amazon Lex returns a `PreconditionFailedException` exception and doesn't publish a new version. If you don't specify a checksum, Amazon Lex publishes the \$LATEST version.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 201
Content-type: application/json

{
  "checksum": "string",
  "conclusionStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "confirmationPrompt": {
    "maxAttempts": number,
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
}
```

```
"createdDate": number,
"description": "string",
"dialogCodeHook": {
  "messageVersion": "string",
  "uri": "string"
},
"followUpPrompt": {
  "prompt": {
    "maxAttempts": number,
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "rejectionStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  }
},
"fulfillmentActivity": {
  "codeHook": {
    "messageVersion": "string",
    "uri": "string"
  },
  "type": "string"
},
"inputContexts": [
  {
    "name": "string"
  }
],
"kendraConfiguration": {
  "kendraIndex": "string",
  "queryFilterString": "string",
```

```

    "role": "string"
  },
  "lastUpdatedDate": number,
  "name": "string",
  "outputContexts": [
    {
      "name": "string",
      "timeToLiveInSeconds": number,
      "turnsToLive": number
    }
  ],
  "parentIntentSignature": "string",
  "rejectionStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupName": number
      }
    ],
    "responseCard": "string"
  },
  "sampleUtterances": [ "string" ],
  "slots": [
    {
      "defaultValueSpec": {
        "defaultValueList": [
          {
            "defaultValue": "string"
          }
        ]
      },
      "description": "string",
      "name": "string",
      "obfuscationSetting": "string",
      "priority": number,
      "responseCard": "string",
      "sampleUtterances": [ "string" ],
      "slotConstraint": "string",
      "slotType": "string",
      "slotTypeVersion": "string",
      "valueElicitationPrompt": {
        "maxAttempts": number,
        "messages": [

```

```
{
  {
    "content": "string",
    "contentType": "string",
    "groupNumber": number
  },
  "responseCard": "string"
},
"version": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 201 response.

The following data is returned in JSON format by the service.

checksum

Checksum of the intent version created.

Type: String

conclusionStatement

After the Lambda function specified in the fulfillmentActivity field fulfills the intent, Amazon Lex conveys this statement to the user.

Type: [Statement](#) object

confirmationPrompt

If defined, the prompt that Amazon Lex uses to confirm the user's intent before fulfilling it.

Type: [Prompt](#) object

createdDate

The date that the intent was created.

Type: Timestamp

description

A description of the intent.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

dialogCodeHook

If defined, Amazon Lex invokes this Lambda function for each user input.

Type: [CodeHook](#) object

followUpPrompt

If defined, Amazon Lex uses this prompt to solicit additional user activity after the intent is fulfilled.

Type: [FollowUpPrompt](#) object

fulfillmentActivity

Describes how the intent is fulfilled.

Type: [FulfillmentActivity](#) object

inputContexts

An array of InputContext objects that lists the contexts that must be active for Amazon Lex to choose the intent in a conversation with the user.

Type: Array of [InputContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 5 items.

kendraConfiguration

Configuration information, if any, for connecting an Amazon Kendra index with the AMAZON.KendraSearchIntent intent.

Type: [KendraConfiguration](#) object

lastUpdatedDate

The date that the intent was updated.

Type: Timestamp

name

The name of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

outputContexts

An array of `OutputContext` objects that lists the contexts that the intent activates when the intent is fulfilled.

Type: Array of [OutputContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

parentIntentSignature

A unique identifier for a built-in intent.

Type: String

rejectionStatement

If the user answers "no" to the question defined in `confirmationPrompt`, Amazon Lex responds with this statement to acknowledge that the intent was canceled.

Type: [Statement](#) object

sampleUtterances

An array of sample utterances configured for the intent.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 1500 items.

Length Constraints: Minimum length of 1. Maximum length of 200.

slots

An array of slot types that defines the information required to fulfill the intent.

Type: Array of [Slot](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

version

The version number assigned to the new version of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

PreconditionFailedException

The checksum of the resource that you are trying to change does not match the checksum in the request. Check the resource's checksum and try again.

HTTP Status Code: 412

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

CreateSlotTypeVersion

Service: Amazon Lex Model Building Service

Creates a new version of a slot type based on the `$LATEST` version of the specified slot type. If the `$LATEST` version of this resource has not changed since the last version that you created, Amazon Lex doesn't create a new version. It returns the last version that you created.

Note

You can update only the `$LATEST` version of a slot type. You can't update the numbered versions that you create with the `CreateSlotTypeVersion` operation.

When you create a version of a slot type, Amazon Lex sets the version to 1. Subsequent versions increment by 1. For more information, see [Versioning](#).

This operation requires permissions for the `lex:CreateSlotTypeVersion` action.

Request Syntax

```
POST /slottypes/name/versions HTTP/1.1
Content-type: application/json

{
  "checksum": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the slot type that you want to create a new version for. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request accepts the following data in JSON format.

checksum

Checksum for the \$LATEST version of the slot type that you want to publish. If you specify a checksum and the \$LATEST version of the slot type has a different checksum, Amazon Lex returns a `PreconditionFailedException` exception and doesn't publish the new version. If you don't specify a checksum, Amazon Lex publishes the \$LATEST version.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 201
Content-type: application/json

{
  "checksum": "string",
  "createdDate": number,
  "description": "string",
  "enumerationValues": [
    {
      "synonyms": [ "string" ],
      "value": "string"
    }
  ],
  "lastUpdatedDate": number,
  "name": "string",
  "parentSlotTypeSignature": "string",
  "slotTypeConfigurations": [
    {
      "regexConfiguration": {
        "pattern": "string"
      }
    }
  ],
  "valueSelectionStrategy": "string",
  "version": "string"
```

```
}
```

Response Elements

If the action is successful, the service sends back an HTTP 201 response.

The following data is returned in JSON format by the service.

[checksum](#)

Checksum of the \$LATEST version of the slot type.

Type: String

[createdDate](#)

The date that the slot type was created.

Type: Timestamp

[description](#)

A description of the slot type.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

[enumerationValues](#)

A list of `EnumerationValue` objects that defines the values that the slot type can take.

Type: Array of [EnumerationValue](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10000 items.

[lastUpdatedDate](#)

The date that the slot type was updated. When you create a resource, the creation date and last update date are the same.

Type: Timestamp

[name](#)

The name of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

parentSlotTypeSignature

The built-in slot type used as the parent of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^((AMAZON\. )_?|[A-Za-z]_?)+`

slotTypeConfigurations

Configuration information that extends the parent built-in slot type.

Type: Array of [SlotTypeConfiguration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

valueSelectionStrategy

The strategy that Amazon Lex uses to determine the value of the slot. For more information, see [PutSlotType](#).

Type: String

Valid Values: ORIGINAL_VALUE | TOP_RESOLUTION

version

The version assigned to the new slot type version.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

PreconditionFailedException

The checksum of the resource that you are trying to change does not match the checksum in the request. Check the resource's checksum and try again.

HTTP Status Code: 412

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteBot

Service: Amazon Lex Model Building Service

Deletes all versions of the bot, including the `$LATEST` version. To delete a specific version of the bot, use the [DeleteBotVersion](#) operation. The `DeleteBot` operation doesn't immediately remove the bot schema. Instead, it is marked for deletion and removed later.

Amazon Lex stores utterances indefinitely for improving the ability of your bot to respond to user inputs. These utterances are not removed when the bot is deleted. To remove the utterances, use the [DeleteUtterances](#) operation.

If a bot has an alias, you can't delete it. Instead, the `DeleteBot` operation returns a `ResourceInUseException` exception that includes a reference to the alias that refers to the bot. To remove the reference to the bot, delete the alias. If you get the same exception again, delete the referring alias until the `DeleteBot` operation is successful.

This operation requires permissions for the `lex:DeleteBot` action.

Request Syntax

```
DELETE /bots/name HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the bot. The name is case sensitive.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

ResourceInUseException

The resource that you are attempting to delete is referred to by another resource. Use this information to remove references to the resource that you are trying to delete.

The body of the exception contains a JSON object that describes the resource.

```
{ "resourceType": BOT | BOTALIAS | BOTCHANNEL | INTENT,  
  "resourceReference": {  
    "name": string, "version": string } }
```

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteBotAlias

Service: Amazon Lex Model Building Service

Deletes an alias for the specified bot.

You can't delete an alias that is used in the association between a bot and a messaging channel. If an alias is used in a channel association, the `DeleteBot` operation returns a `ResourceInUseException` exception that includes a reference to the channel association that refers to the bot. You can remove the reference to the alias by deleting the channel association. If you get the same exception again, delete the referring association until the `DeleteBotAlias` operation is successful.

Request Syntax

```
DELETE /bots/botName/aliases/name HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

botName

The name of the bot that the alias points to.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

name

The name of the alias to delete. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

ResourceInUseException

The resource that you are attempting to delete is referred to by another resource. Use this information to remove references to the resource that you are trying to delete.

The body of the exception contains a JSON object that describes the resource.

```
{ "resourceType": BOT | BOTALIAS | BOTCHANNEL | INTENT,  
  "resourceReference": {  
    "name": string, "version": string } }
```

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteBotChannelAssociation

Service: Amazon Lex Model Building Service

Deletes the association between an Amazon Lex bot and a messaging platform.

This operation requires permission for the `lex:DeleteBotChannelAssociation` action.

Request Syntax

```
DELETE /bots/botName/aliases/aliasName/channels/name HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

aliasName

An alias that points to the specific version of the Amazon Lex bot to which this association is being made.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

botName

The name of the Amazon Lex bot.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

name

The name of the association. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteBotVersion

Service: Amazon Lex Model Building Service

Deletes a specific version of a bot. To delete all versions of a bot, use the [DeleteBot](#) operation.

This operation requires permissions for the `lex:DeleteBotVersion` action.

Request Syntax

```
DELETE /bots/name/versions/version HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the bot.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

version

The version of the bot to delete. You cannot delete the `$LATEST` version of the bot. To delete the `$LATEST` version, use the [DeleteBot](#) operation.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

ResourceInUseException

The resource that you are attempting to delete is referred to by another resource. Use this information to remove references to the resource that you are trying to delete.

The body of the exception contains a JSON object that describes the resource.

```
{ "resourceType": BOT | BOTALIAS | BOTCHANNEL | INTENT,  
  "resourceReference": {
```

```
"name": string, "version": string } }
```

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteIntent

Service: Amazon Lex Model Building Service

Deletes all versions of the intent, including the \$LATEST version. To delete a specific version of the intent, use the [DeleteIntentVersion](#) operation.

You can delete a version of an intent only if it is not referenced. To delete an intent that is referred to in one or more bots (see [Amazon Lex: How It Works](#)), you must remove those references first.

Note

If you get the `ResourceInUseException` exception, it provides an example reference that shows where the intent is referenced. To remove the reference to the intent, either update the bot or delete it. If you get the same exception when you attempt to delete the intent again, repeat until the intent has no references and the call to `DeleteIntent` is successful.

This operation requires permission for the `lex:DeleteIntent` action.

Request Syntax

```
DELETE /intents/name HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the intent. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

ResourceInUseException

The resource that you are attempting to delete is referred to by another resource. Use this information to remove references to the resource that you are trying to delete.

The body of the exception contains a JSON object that describes the resource.

```
{ "resourceType": BOT | BOTALIAS | BOTCHANNEL | INTENT,  
  "resourceReference": {  
    "name": string, "version": string } }
```

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteIntentVersion

Service: Amazon Lex Model Building Service

Deletes a specific version of an intent. To delete all versions of a intent, use the [DeleteIntent](#) operation.

This operation requires permissions for the `lex:DeleteIntentVersion` action.

Request Syntax

```
DELETE /intents/name/versions/version HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the intent.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

version

The version of the intent to delete. You cannot delete the `$LATEST` version of the intent. To delete the `$LATEST` version, use the [DeleteIntent](#) operation.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

ResourceInUseException

The resource that you are attempting to delete is referred to by another resource. Use this information to remove references to the resource that you are trying to delete.

The body of the exception contains a JSON object that describes the resource.

```
{ "resourceType": BOT | BOTALIAS | BOTCHANNEL | INTENT,  
  "resourceReference": {  
    "name": string, "version": string } }
```

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteSlotType

Service: Amazon Lex Model Building Service

Deletes all versions of the slot type, including the `$LATEST` version. To delete a specific version of the slot type, use the [DeleteSlotTypeVersion](#) operation.

You can delete a version of a slot type only if it is not referenced. To delete a slot type that is referred to in one or more intents, you must remove those references first.

Note

If you get the `ResourceInUseException` exception, the exception provides an example reference that shows the intent where the slot type is referenced. To remove the reference to the slot type, either update the intent or delete it. If you get the same exception when you attempt to delete the slot type again, repeat until the slot type has no references and the `DeleteSlotType` call is successful.

This operation requires permission for the `lex:DeleteSlotType` action.

Request Syntax

```
DELETE /slottypes/name HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the slot type. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

ResourceInUseException

The resource that you are attempting to delete is referred to by another resource. Use this information to remove references to the resource that you are trying to delete.

The body of the exception contains a JSON object that describes the resource.

```
{ "resourceType": BOT | BOTALIAS | BOTCHANNEL | INTENT,  
  "resourceReference": {  
    "name": string, "version": string } }
```

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteSlotTypeVersion

Service: Amazon Lex Model Building Service

Deletes a specific version of a slot type. To delete all versions of a slot type, use the [DeleteSlotType](#) operation.

This operation requires permissions for the `lex:DeleteSlotTypeVersion` action.

Request Syntax

```
DELETE /slottypes/name/version/version HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the slot type.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

version

The version of the slot type to delete. You cannot delete the `$LATEST` version of the slot type. To delete the `$LATEST` version, use the [DeleteSlotType](#) operation.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

ResourceInUseException

The resource that you are attempting to delete is referred to by another resource. Use this information to remove references to the resource that you are trying to delete.

The body of the exception contains a JSON object that describes the resource.

```
{ "resourceType": BOT | BOTALIAS | BOTCHANNEL | INTENT,  
  "resourceReference": {  
    "name": string, "version": string } }
```

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteUtterances

Service: Amazon Lex Model Building Service

Deletes stored utterances.

Amazon Lex stores the utterances that users send to your bot. Utterances are stored for 15 days for use with the [GetUtterancesView](#) operation, and then stored indefinitely for use in improving the ability of your bot to respond to user input.

Use the DeleteUtterances operation to manually delete stored utterances for a specific user. When you use the DeleteUtterances operation, utterances stored for improving your bot's ability to respond to user input are deleted immediately. Utterances stored for use with the GetUtterancesView operation are deleted after 15 days.

This operation requires permissions for the `lex:DeleteUtterances` action.

Request Syntax

```
DELETE /bots/botName/utterances/userId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

botName

The name of the bot that stored the utterances.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

userId

The unique identifier for the user that made the utterances. This is the user ID that was sent in the [PostContent](#) or [PostText](#) operation request that contained the utterance.

Length Constraints: Minimum length of 2. Maximum length of 100.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetBot

Service: Amazon Lex Model Building Service

Returns metadata information for a specific bot. You must provide the bot name and the bot version or alias.

This operation requires permissions for the `lex:GetBot` action.

Request Syntax

```
GET /bots/name/versions/versionoralias HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the bot. The name is case sensitive.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

versionoralias

The version or alias of the bot.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "abortStatement": {
    "messages": [
```

```

    {
      "content": "string",
      "contentType": "string",
      "groupNumber": number
    }
  ],
  "responseCard": "string"
},
"checksum": "string",
"childDirected": boolean,
"clarificationPrompt": {
  "maxAttempts": number,
  "messages": [
    {
      "content": "string",
      "contentType": "string",
      "groupNumber": number
    }
  ],
  "responseCard": "string"
},
"createdDate": number,
"description": "string",
"detectSentiment": boolean,
"enableModelImprovements": boolean,
"failureReason": "string",
"idleSessionTTLInSeconds": number,
"intents": [
  {
    "intentName": "string",
    "intentVersion": "string"
  }
],
"lastUpdatedDate": number,
"locale": "string",
"name": "string",
"nluIntentConfidenceThreshold": number,
"status": "string",
"version": "string",
"voiceId": "string"
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

abortStatement

The message that Amazon Lex returns when the user elects to end the conversation without completing it. For more information, see [PutBot](#).

Type: [Statement](#) object

checksum

Checksum of the bot used to identify a specific revision of the bot's \$LATEST version.

Type: String

childDirected

For each Amazon Lex bot created with the Amazon Lex Model Building Service, you must specify whether your use of Amazon Lex is related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to the Children's Online Privacy Protection Act (COPPA) by specifying true or false in the `childDirected` field. By specifying true in the `childDirected` field, you confirm that your use of Amazon Lex **is** related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA. By specifying false in the `childDirected` field, you confirm that your use of Amazon Lex **is not** related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA. You may not specify a default value for the `childDirected` field that does not accurately reflect whether your use of Amazon Lex is related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA.

If your use of Amazon Lex relates to a website, program, or other application that is directed in whole or in part, to children under age 13, you must obtain any required verifiable parental consent under COPPA. For information regarding the use of Amazon Lex in connection with websites, programs, or other applications that are directed or targeted, in whole or in part, to children under age 13, see the [Amazon Lex FAQ](#).

Type: Boolean

clarificationPrompt

The message Amazon Lex uses when it doesn't understand the user's request. For more information, see [PutBot](#).

Type: [Prompt](#) object

createdDate

The date that the bot was created.

Type: Timestamp

description

A description of the bot.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

detectSentiment

Indicates whether user utterances should be sent to Amazon Comprehend for sentiment analysis.

Type: Boolean

enableModelImprovements

Indicates whether the bot uses accuracy improvements. `true` indicates that the bot is using the improvements, otherwise, `false`.

Type: Boolean

failureReason

If status is `FAILED`, Amazon Lex explains why it failed to build the bot.

Type: String

idleSessionTTLInSeconds

The maximum time in seconds that Amazon Lex retains the data gathered in a conversation. For more information, see [PutBot](#).

Type: Integer

Valid Range: Minimum value of 60. Maximum value of 86400.

intents

An array of intent objects. For more information, see [PutBot](#).

Type: Array of [Intent](#) objects

lastUpdatedDate

The date that the bot was updated. When you create a resource, the creation date and last updated date are the same.

Type: Timestamp

locale

The target locale for the bot.

Type: String

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

name

The name of the bot.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z]_?)+\$

nluIntentConfidenceThreshold

The score that determines where Amazon Lex inserts the `AMAZON.FallbackIntent`, `AMAZON.KendraSearchIntent`, or both when returning alternative intents in a [PostContent](#) or [PostText](#) response. `AMAZON.FallbackIntent` is inserted if the confidence score for all intents is below this value. `AMAZON.KendraSearchIntent` is only inserted if it is configured for the bot.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 1.

status

The status of the bot.

When the status is `BUILDING` Amazon Lex is building the bot for testing and use.

If the status of the bot is `READY_BASIC_TESTING`, you can test the bot using the exact utterances specified in the bot's intents. When the bot is ready for full testing or to run, the status is `READY`.

If there was a problem with building the bot, the status is `FAILED` and the `failureReason` field explains why the bot did not build.

If the bot was saved but not built, the status is `NOT_BUILT`.

Type: String

Valid Values: `BUILDING` | `READY` | `READY_BASIC_TESTING` | `FAILED` | `NOT_BUILT`

version

The version of the bot. For a new bot, the version is always `$LATEST`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

voiceId

The Amazon Polly voice ID that Amazon Lex uses for voice interaction with the user. For more information, see [PutBot](#).

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetBotAlias

Service: Amazon Lex Model Building Service

Returns information about an Amazon Lex bot alias. For more information about aliases, see [Versioning and Aliases](#).

This operation requires permissions for the `lex:GetBotAlias` action.

Request Syntax

```
GET /bots/botName/aliases/name HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

botName

The name of the bot.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

name

The name of the bot alias. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Content-type: application/json

```
{
  "botName": "string",
  "botVersion": "string",
  "checksum": "string",
  "conversationLogs": {
    "iamRoleArn": "string",
    "logSettings": [
      {
        "destination": "string",
        "kmsKeyArn": "string",
        "logType": "string",
        "resourceArn": "string",
        "resourcePrefix": "string"
      }
    ]
  },
  "createdDate": number,
  "description": "string",
  "lastUpdatedDate": number,
  "name": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

botName

The name of the bot that the alias points to.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z]_?)+\$

botVersion

The version of the bot that the alias points to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

checksum

Checksum of the bot alias.

Type: String

conversationLogs

The settings that determine how Amazon Lex uses conversation logs for the alias.

Type: [ConversationLogsResponse](#) object

createdDate

The date that the bot alias was created.

Type: Timestamp

description

A description of the bot alias.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

lastUpdatedDate

The date that the bot alias was updated. When you create a resource, the creation date and the last updated date are the same.

Type: Timestamp

name

The name of the bot alias.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

GetBotAliases

Service: Amazon Lex Model Building Service

Returns a list of aliases for a specified Amazon Lex bot.

This operation requires permissions for the `lex:GetBotAliases` action.

Request Syntax

```
GET /bots/botName/aliases/?  
maxResults=maxResults&nameContains=nameContains&nextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

botName

The name of the bot.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

maxResults

The maximum number of aliases to return in the response. The default is 50. .

Valid Range: Minimum value of 1. Maximum value of 50.

nameContains

Substring to match in bot alias names. An alias will be returned if any part of its name matches the substring. For example, "xyz" matches both "xyzabc" and "abcxyz."

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

nextToken

A pagination token for fetching the next page of aliases. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of aliases, specify the pagination token in the next request.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "BotAliases": [
    {
      "botName": "string",
      "botVersion": "string",
      "checksum": "string",
      "conversationLogs": {
        "iamRoleArn": "string",
        "logSettings": [
          {
            "destination": "string",
            "kmsKeyArn": "string",
            "logType": "string",
            "resourceArn": "string",
            "resourcePrefix": "string"
          }
        ]
      },
      "createdDate": number,
      "description": "string",
      "lastUpdatedDate": number,
      "name": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

BotAliases

An array of `BotAliasMetadata` objects, each describing a bot alias.

Type: Array of [BotAliasMetadata](#) objects

nextToken

A pagination token for fetching next page of aliases. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of aliases, specify the pagination token in the next request.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetBotChannelAssociation

Service: Amazon Lex Model Building Service

Returns information about the association between an Amazon Lex bot and a messaging platform.

This operation requires permissions for the `lex:GetBotChannelAssociation` action.

Request Syntax

```
GET /bots/botName/aliases/aliasName/channels/name HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

aliasName

An alias pointing to the specific version of the Amazon Lex bot to which this association is being made.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

botName

The name of the Amazon Lex bot.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

name

The name of the association between the bot and the channel. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "botAlias": "string",
  "botConfiguration": {
    "string" : "string"
  },
  "botName": "string",
  "createdDate": number,
  "description": "string",
  "failureReason": "string",
  "name": "string",
  "status": "string",
  "type": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

botAlias

An alias pointing to the specific version of the Amazon Lex bot to which this association is being made.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: ^([A-Za-z_?])+ \$

botConfiguration

Provides information that the messaging platform needs to communicate with the Amazon Lex bot.

Type: String to string map

Map Entries: Maximum number of 10 items.

botName

The name of the Amazon Lex bot.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

createdDate

The date that the association between the bot and the channel was created.

Type: Timestamp

description

A description of the association between the bot and the channel.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

failureReason

If status is FAILED, Amazon Lex provides the reason that it failed to create the association.

Type: String

name

The name of the association between the bot and the channel.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

status

The status of the bot channel.

- **CREATED** - The channel has been created and is ready for use.
- **IN_PROGRESS** - Channel creation is in progress.
- **FAILED** - There was an error creating the channel. For information about the reason for the failure, see the `failureReason` field.

Type: String

Valid Values: `IN_PROGRESS` | `CREATED` | `FAILED`

type

The type of the messaging platform.

Type: String

Valid Values: `Facebook` | `Slack` | `Twilio-Sms` | `Kik`

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetBotChannelAssociations

Service: Amazon Lex Model Building Service

Returns a list of all of the channels associated with the specified bot.

The GetBotChannelAssociations operation requires permissions for the `lex:GetBotChannelAssociations` action.

Request Syntax

```
GET /bots/botName/aliases/aliasName/channels/?  
maxResults=maxResults&nameContains=nameContains&nextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

aliasName

An alias pointing to the specific version of the Amazon Lex bot to which this association is being made.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^(-|^[A-Za-z]_?)+$`

Required: Yes

botName

The name of the Amazon Lex bot in the association.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z]_?)+$`

Required: Yes

maxResults

The maximum number of associations to return in the response. The default is 50.

Valid Range: Minimum value of 1. Maximum value of 50.

nameContains

Substring to match in channel association names. An association will be returned if any part of its name matches the substring. For example, "xyz" matches both "xyzabc" and "abcxyz." To return all bot channel associations, use a hyphen ("-") as the nameContains parameter.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

nextToken

A pagination token for fetching the next page of associations. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of associations, specify the pagination token in the next request.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "botChannelAssociations": [
    {
      "botAlias": "string",
      "botConfiguration": {
        "string" : "string"
      },
      "botName": "string",
      "createdDate": number,
      "description": "string",
      "failureReason": "string",
      "name": "string",
      "status": "string",
      "type": "string"
    }
  ],
  "nextToken": "string"
```

```
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[botChannelAssociations](#)

An array of objects, one for each association, that provides information about the Amazon Lex bot and its association with the channel.

Type: Array of [BotChannelAssociation](#) objects

[nextToken](#)

A pagination token that fetches the next page of associations. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of associations, specify the pagination token in the next request.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetBots

Service: Amazon Lex Model Building Service

Returns bot information as follows:

- If you provide the `nameContains` field, the response includes information for the \$LATEST version of all bots whose name contains the specified string.
- If you don't specify the `nameContains` field, the operation returns information about the \$LATEST version of all of your bots.

This operation requires permission for the `lex:GetBots` action.

Request Syntax

```
GET /bots/?maxResults=maxResults&nameContains=nameContains&nextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of bots to return in the response that the request will return. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

nameContains

Substring to match in bot names. A bot will be returned if any part of its name matches the substring. For example, "xyz" matches both "xyzabc" and "abcxyz."

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^([A-Za-z]_?)+$`

nextToken

A pagination token that fetches the next page of bots. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of bots, specify the pagination token in the next request.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "bots": [
    {
      "createdDate": number,
      "description": "string",
      "lastUpdatedDate": number,
      "name": "string",
      "status": "string",
      "version": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

bots

An array of botMetadata objects, with one entry for each bot.

Type: Array of [BotMetadata](#) objects

nextToken

If the response is truncated, it includes a pagination token that you can specify in your next request to fetch the next page of bots.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

GetBotVersions

Service: Amazon Lex Model Building Service

Gets information about all of the versions of a bot.

The GetBotVersions operation returns a BotMetadata object for each version of a bot. For example, if a bot has three numbered versions, the GetBotVersions operation returns four BotMetadata objects in the response, one for each numbered version and one for the \$LATEST version.

The GetBotVersions operation always returns at least one version, the \$LATEST version.

This operation requires permissions for the `lex:GetBotVersions` action.

Request Syntax

```
GET /bots/name/versions/?maxResults=maxResults&nextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of bot versions to return in the response. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

name

The name of the bot for which versions should be returned.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

nextToken

A pagination token for fetching the next page of bot versions. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of versions, specify the pagination token in the next request.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "bots": [
    {
      "createdDate": number,
      "description": "string",
      "lastUpdatedDate": number,
      "name": "string",
      "status": "string",
      "version": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

bots

An array of BotMetadata objects, one for each numbered version of the bot plus one for the \$LATEST version.

Type: Array of [BotMetadata](#) objects

nextToken

A pagination token for fetching the next page of bot versions. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of versions, specify the pagination token in the next request.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

GetBuiltinIntent

Service: Amazon Lex Model Building Service

Returns information about a built-in intent.

This operation requires permission for the `lex:GetBuiltinIntent` action.

Request Syntax

```
GET /builtins/intents/signature HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

signature

The unique identifier for a built-in intent. To find the signature for an intent, see [Standard Built-in Intents](#) in the *Alexa Skills Kit*.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "signature": "string",
  "slots": [
    {
      "name": "string"
    }
  ],
  "supportedLocales": [ "string" ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

signature

The unique identifier for a built-in intent.

Type: String

slots

An array of `BuiltinIntentSlot` objects, one entry for each slot type in the intent.

Type: Array of [BuiltinIntentSlot](#) objects

supportedLocales

A list of locales that the intent supports.

Type: Array of strings

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetBuiltinIntents

Service: Amazon Lex Model Building Service

Gets a list of built-in intents that meet the specified criteria.

This operation requires permission for the `lex:GetBuiltinIntents` action.

Request Syntax

```
GET /builtins/intents/?  
locale=locale&maxResults=maxResults&nextToken=nextToken&signatureContains=signatureContains  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

locale

A list of locales that the intent supports.

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

maxResults

The maximum number of intents to return in the response. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

nextToken

A pagination token that fetches the next page of intents. If this API call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of intents, use the pagination token in the next request.

signatureContains

Substring to match in built-in intent signatures. An intent will be returned if any part of its signature matches the substring. For example, "xyz" matches both "xyzabc" and "abcxyz." To find the signature for an intent, see [Standard Built-in Intents](#) in the *Alexa Skills Kit*.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "intents": [
    {
      "signature": "string",
      "supportedLocales": [ "string" ]
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

intents

An array of `BuiltinIntentMetadata` objects, one for each intent in the response.

Type: Array of [BuiltinIntentMetadata](#) objects

nextToken

A pagination token that fetches the next page of intents. If the response to this API call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of intents, specify the pagination token in the next request.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetBuiltinSlotTypes

Service: Amazon Lex Model Building Service

Gets a list of built-in slot types that meet the specified criteria.

For a list of built-in slot types, see [Slot Type Reference](#) in the *Alexa Skills Kit*.

This operation requires permission for the `lex:GetBuiltinSlotTypes` action.

Request Syntax

```
GET /builtins/slottypes/?  
locale=locale&maxResults=maxResults&nextToken=nextToken&signatureContains=signatureContains  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

locale

A list of locales that the slot type supports.

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

maxResults

The maximum number of slot types to return in the response. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

nextToken

A pagination token that fetches the next page of slot types. If the response to this API call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of slot types, specify the pagination token in the next request.

signatureContains

Substring to match in built-in slot type signatures. A slot type will be returned if any part of its signature matches the substring. For example, "xyz" matches both "xyzabc" and "abcxyz."

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "slotTypes": [
    {
      "signature": "string",
      "supportedLocales": [ "string" ]
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[nextToken](#)

If the response is truncated, the response includes a pagination token that you can use in your next request to fetch the next page of slot types.

Type: String

[slotTypes](#)

An array of `BuiltInSlotTypeMetadata` objects, one entry for each slot type returned.

Type: Array of [BuiltinSlotTypeMetadata](#) objects

Errors

`BadRequestException`

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetExport

Service: Amazon Lex Model Building Service

Exports the contents of a Amazon Lex resource in a specified format.

Request Syntax

```
GET /exports/?exportType=exportType&name=name&resourceType=resourceType&version=version
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

exportType

The format of the exported data.

Valid Values: ALEXA_SKILLS_KIT | LEX

Required: Yes

name

The name of the bot to export.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: [a-zA-Z_]+

Required: Yes

resourceType

The type of resource to export.

Valid Values: BOT | INTENT | SLOT_TYPE

Required: Yes

version

The version of the bot to export.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9] +

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "exportStatus": "string",
  "exportType": "string",
  "failureReason": "string",
  "name": "string",
  "resourceType": "string",
  "url": "string",
  "version": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

exportStatus

The status of the export.

- IN_PROGRESS - The export is in progress.
- READY - The export is complete.
- FAILED - The export could not be completed.

Type: String

Valid Values: IN_PROGRESS | READY | FAILED

exportType

The format of the exported data.

Type: String

Valid Values: ALEXA_SKILLS_KIT | LEX

failureReason

If status is FAILED, Amazon Lex provides the reason that it failed to export the resource.

Type: String

name

The name of the bot being exported.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: [a-zA-Z_]+

resourceType

The type of the exported resource.

Type: String

Valid Values: BOT | INTENT | SLOT_TYPE

url

An S3 pre-signed URL that provides the location of the exported resource. The exported resource is a ZIP archive that contains the exported resource in JSON format. The structure of the archive may change. Your code should not rely on the archive structure.

Type: String

version

The version of the bot being exported.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9]+

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetImport

Service: Amazon Lex Model Building Service

Gets information about an import job started with the `StartImport` operation.

Request Syntax

```
GET /imports/importId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

importId

The identifier of the import job information to return.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "createdDate": number,
  "failureReason": [ "string" ],
  "importId": "string",
  "importStatus": "string",
  "mergeStrategy": "string",
  "name": "string",
  "resourceType": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

createdDate

A timestamp for the date and time that the import job was created.

Type: Timestamp

failureReason

A string that describes why an import job failed to complete.

Type: Array of strings

importId

The identifier for the specific import job.

Type: String

importStatus

The status of the import job. If the status is FAILED, you can get the reason for the failure from the `failureReason` field.

Type: String

Valid Values: IN_PROGRESS | COMPLETE | FAILED

mergeStrategy

The action taken when there was a conflict between an existing resource and a resource in the import file.

Type: String

Valid Values: OVERWRITE_LATEST | FAIL_ON_CONFLICT

name

The name given to the import job.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `[a-zA-Z_]+`

resourceType

The type of resource imported.

Type: String

Valid Values: BOT | INTENT | SLOT_TYPE

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetIntent

Service: Amazon Lex Model Building Service

Returns information about an intent. In addition to the intent name, you must specify the intent version.

This operation requires permissions to perform the `lex:GetIntent` action.

Request Syntax

```
GET /intents/name/versions/version HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the intent. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

version

The version of the intent.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200  
Content-type: application/json
```

```
{
  "checksum": "string",
  "conclusionStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "confirmationPrompt": {
    "maxAttempts": number,
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "createdAt": number,
  "description": "string",
  "dialogCodeHook": {
    "messageVersion": "string",
    "uri": "string"
  },
  "followUpPrompt": {
    "prompt": {
      "maxAttempts": number,
      "messages": [
        {
          "content": "string",
          "contentType": "string",
          "groupNumber": number
        }
      ]
    },
    "responseCard": "string"
  },
  "rejectionStatement": {
    "messages": [
```



```
        {
            "content": "string",
            "contentType": "string",
            "groupNumber": number
        }
    ],
    "responseCard": "string"
}
},
"fulfillmentActivity": {
    "codeHook": {
        "messageVersion": "string",
        "uri": "string"
    },
    "type": "string"
},
"inputContexts": [
    {
        "name": "string"
    }
],
"kendraConfiguration": {
    "kendraIndex": "string",
    "queryFilterString": "string",
    "role": "string"
},
"lastUpdatedDate": number,
"name": "string",
"outputContexts": [
    {
        "name": "string",
        "timeToLiveInSeconds": number,
        "turnsToLive": number
    }
],
"parentIntentSignature": "string",
"rejectionStatement": {
    "messages": [
        {
            "content": "string",
            "contentType": "string",
            "groupNumber": number
        }
    ]
},
],
```

```

    "responseCard": "string"
  },
  "sampleUtterances": [ "string" ],
  "slots": [
    {
      "defaultValueSpec": {
        "defaultValueList": [
          {
            "defaultValue": "string"
          }
        ]
      },
      "description": "string",
      "name": "string",
      "obfuscationSetting": "string",
      "priority": number,
      "responseCard": "string",
      "sampleUtterances": [ "string" ],
      "slotConstraint": "string",
      "slotType": "string",
      "slotTypeVersion": "string",
      "valueElicitationPrompt": {
        "maxAttempts": number,
        "messages": [
          {
            "content": "string",
            "contentType": "string",
            "groupNumber": number
          }
        ],
        "responseCard": "string"
      }
    }
  ],
  "version": "string"
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

checksum

Checksum of the intent.

Type: String

conclusionStatement

After the Lambda function specified in the `fulfillmentActivity` element fulfills the intent, Amazon Lex conveys this statement to the user.

Type: [Statement](#) object

confirmationPrompt

If defined in the bot, Amazon Lex uses prompt to confirm the intent before fulfilling the user's request. For more information, see [PutIntent](#).

Type: [Prompt](#) object

createdDate

The date that the intent was created.

Type: Timestamp

description

A description of the intent.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

dialogCodeHook

If defined in the bot, Amazon Amazon Lex invokes this Lambda function for each user input. For more information, see [PutIntent](#).

Type: [CodeHook](#) object

followUpPrompt

If defined in the bot, Amazon Lex uses this prompt to solicit additional user activity after the intent is fulfilled. For more information, see [PutIntent](#).

Type: [FollowUpPrompt](#) object

fulfillmentActivity

Describes how the intent is fulfilled. For more information, see [PutIntent](#).

Type: [FulfillmentActivity](#) object

inputContexts

An array of InputContext objects that lists the contexts that must be active for Amazon Lex to choose the intent in a conversation with the user.

Type: Array of [InputContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 5 items.

kendraConfiguration

Configuration information, if any, to connect to an Amazon Kendra index with the AMAZON.KendraSearchIntent intent.

Type: [KendraConfiguration](#) object

lastUpdatedDate

The date that the intent was updated. When you create a resource, the creation date and the last updated date are the same.

Type: Timestamp

name

The name of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

outputContexts

An array of OutputContext objects that lists the contexts that the intent activates when the intent is fulfilled.

Type: Array of [OutputContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

parentIntentSignature

A unique identifier for a built-in intent.

Type: String

rejectionStatement

If the user answers "no" to the question defined in `confirmationPrompt`, Amazon Lex responds with this statement to acknowledge that the intent was canceled.

Type: [Statement](#) object

sampleUtterances

An array of sample utterances configured for the intent.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 1500 items.

Length Constraints: Minimum length of 1. Maximum length of 200.

slots

An array of intent slots configured for the intent.

Type: Array of [Slot](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

version

The version of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetIntentents

Service: Amazon Lex Model Building Service

Returns intent information as follows:

- If you specify the `nameContains` field, returns the \$LATEST version of all intents that contain the specified string.
- If you don't specify the `nameContains` field, returns information about the \$LATEST version of all intents.

The operation requires permission for the `lex:GetIntentents` action.

Request Syntax

```
GET /intentents/?maxResults=maxResults&nameContains=nameContains&nextToken=nextToken
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of intents to return in the response. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

nameContains

Substring to match in intent names. An intent will be returned if any part of its name matches the substring. For example, "xyz" matches both "xyzabc" and "abcxyz."

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z_?])+`

nextToken

A pagination token that fetches the next page of intents. If the response to this API call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of intents, specify the pagination token in the next request.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "intents": [
    {
      "createdDate": number,
      "description": "string",
      "lastUpdatedDate": number,
      "name": "string",
      "version": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

intents

An array of Intent objects. For more information, see [PutBot](#).

Type: Array of [IntentMetadata](#) objects

nextToken

If the response is truncated, the response includes a pagination token that you can specify in your next request to fetch the next page of intents.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

GetIntentVersions

Service: Amazon Lex Model Building Service

Gets information about all of the versions of an intent.

The `GetIntentVersions` operation returns an `IntentMetadata` object for each version of an intent. For example, if an intent has three numbered versions, the `GetIntentVersions` operation returns four `IntentMetadata` objects in the response, one for each numbered version and one for the `$LATEST` version.

The `GetIntentVersions` operation always returns at least one version, the `$LATEST` version.

This operation requires permissions for the `lex:GetIntentVersions` action.

Request Syntax

```
GET /intents/name/versions/?maxResults=maxResults&nextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of intent versions to return in the response. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

name

The name of the intent for which versions should be returned.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

nextToken

A pagination token for fetching the next page of intent versions. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of versions, specify the pagination token in the next request.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "intents": [
    {
      "createdDate": number,
      "description": "string",
      "lastUpdatedDate": number,
      "name": "string",
      "version": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

intents

An array of IntentMetadata objects, one for each numbered version of the intent plus one for the \$LATEST version.

Type: Array of [IntentMetadata](#) objects

nextToken

A pagination token for fetching the next page of intent versions. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of versions, specify the pagination token in the next request.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

GetMigration

Service: Amazon Lex Model Building Service

Provides details about an ongoing or complete migration from an Amazon Lex V1 bot to an Amazon Lex V2 bot. Use this operation to view the migration alerts and warnings related to the migration.

Request Syntax

```
GET /migrations/migrationId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

migrationId

The unique identifier of the migration to view. The migrationID is returned by the [StartMigration](#) operation.

Length Constraints: Fixed length of 10.

Pattern: `^[0-9a-zA-Z]+$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "alerts": [
    {
      "details": [ "string" ],
      "message": "string",
      "referenceURLs": [ "string" ],
      "type": "string"
    }
  ]
}
```

```
    }  
  ],  
  "migrationId": "string",  
  "migrationStatus": "string",  
  "migrationStrategy": "string",  
  "migrationTimestamp": number,  
  "v1BotLocale": "string",  
  "v1BotName": "string",  
  "v1BotVersion": "string",  
  "v2BotId": "string",  
  "v2BotRole": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

alerts

A list of alerts and warnings that indicate issues with the migration for the Amazon Lex V1 bot to Amazon Lex V2. You receive a warning when an Amazon Lex V1 feature has a different implementation in Amazon Lex V2.

For more information, see [Migrating a bot](#) in the *Amazon Lex V2 developer guide*.

Type: Array of [MigrationAlert](#) objects

migrationId

The unique identifier of the migration. This is the same as the identifier used when calling the `GetMigration` operation.

Type: String

Length Constraints: Fixed length of 10.

Pattern: `^[0-9a-zA-Z]+$`

migrationStatus

Indicates the status of the migration. When the status is `COMPLETE` the migration is finished and the bot is available in Amazon Lex V2. There may be alerts and warnings that need to be resolved to complete the migration.

Type: String

Valid Values: IN_PROGRESS | COMPLETED | FAILED

migrationStrategy

The strategy used to conduct the migration.

- CREATE_NEW - Creates a new Amazon Lex V2 bot and migrates the Amazon Lex V1 bot to the new bot.
- UPDATE_EXISTING - Overwrites the existing Amazon Lex V2 bot metadata and the locale being migrated. It doesn't change any other locales in the Amazon Lex V2 bot. If the locale doesn't exist, a new locale is created in the Amazon Lex V2 bot.

Type: String

Valid Values: CREATE_NEW | UPDATE_EXISTING

migrationTimestamp

The date and time that the migration started.

Type: Timestamp

v1BotLocale

The locale of the Amazon Lex V1 bot migrated to Amazon Lex V2.

Type: String

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

v1BotName

The name of the Amazon Lex V1 bot migrated to Amazon Lex V2.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z]_?)+\$

v1BotVersion

The version of the Amazon Lex V1 bot migrated to Amazon Lex V2.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

[v2BotId](#)

The unique identifier of the Amazon Lex V2 bot that the Amazon Lex V1 is being migrated to.

Type: String

Length Constraints: Fixed length of 10.

Pattern: `^[0-9a-zA-Z]+$`

[v2BotRole](#)

The IAM role that Amazon Lex uses to run the Amazon Lex V2 bot.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `^arn:[\w\-\-]+:iam::[\d]{12}:role/.+$`

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetMigrations

Service: Amazon Lex Model Building Service

Gets a list of migrations between Amazon Lex V1 and Amazon Lex V2.

Request Syntax

```
GET /migrations?  
maxResults=maxResults&migrationStatusEquals=migrationStatusEquals&nextToken=nextToken&sortByAttribute=sortByAttribute  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of migrations to return in the response. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

migrationStatusEquals

Filters the list to contain only migrations in the specified state.

Valid Values: IN_PROGRESS | COMPLETED | FAILED

nextToken

A pagination token that fetches the next page of migrations. If the response to this operation is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of migrations, specify the pagination token in the request.

sortByAttribute

The field to sort the list of migrations by. You can sort by the Amazon Lex V1 bot name or the date and time that the migration was started.

Valid Values: V1_BOT_NAME | MIGRATION_DATE_TIME

sortByOrder

The order to sort the list.

Valid Values: ASCENDING | DESCENDING

v1BotNameContains

Filters the list to contain only bots whose name contains the specified string. The string is matched anywhere in the bot name.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^([A-Za-z]_?)+$`

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "migrationSummaries": [
    {
      "migrationId": "string",
      "migrationStatus": "string",
      "migrationStrategy": "string",
      "migrationTimestamp": number,
      "v1BotLocale": "string",
      "v1BotName": "string",
      "v1BotVersion": "string",
      "v2BotId": "string",
      "v2BotRole": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[migrationSummaries](#)

An array of summaries for migrations from Amazon Lex V1 to Amazon Lex V2. To see details of the migration, use the `migrationId` from the summary in a call to the [GetMigration](#) operation.

Type: Array of [MigrationSummary](#) objects

[nextToken](#)

If the response is truncated, it includes a pagination token that you can specify in your next request to fetch the next page of migrations.

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetSlotType

Service: Amazon Lex Model Building Service

Returns information about a specific version of a slot type. In addition to specifying the slot type name, you must specify the slot type version.

This operation requires permissions for the `lex:GetSlotType` action.

Request Syntax

```
GET /slottypes/name/versions/version HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the slot type. The name is case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

version

The version of the slot type.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200  
Content-type: application/json
```



```
{
  "checksum": "string",
  "createdDate": number,
  "description": "string",
  "enumerationValues": [
    {
      "synonyms": [ "string" ],
      "value": "string"
    }
  ],
  "lastUpdatedDate": number,
  "name": "string",
  "parentSlotTypeSignature": "string",
  "slotTypeConfigurations": [
    {
      "regexConfiguration": {
        "pattern": "string"
      }
    }
  ],
  "valueSelectionStrategy": "string",
  "version": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

checksum

Checksum of the \$LATEST version of the slot type.

Type: String

createdDate

The date that the slot type was created.

Type: Timestamp

description

A description of the slot type.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

enumerationValues

A list of `EnumerationValue` objects that defines the values that the slot type can take.

Type: Array of [EnumerationValue](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10000 items.

lastUpdatedDate

The date that the slot type was updated. When you create a resource, the creation date and last update date are the same.

Type: Timestamp

name

The name of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

parentSlotTypeSignature

The built-in slot type used as a parent for the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^( (AMAZON\.)_? | [A-Za-z_?]+ )`

slotTypeConfigurations

Configuration information that extends the parent built-in slot type.

Type: Array of [SlotTypeConfiguration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

valueSelectionStrategy

The strategy that Amazon Lex uses to determine the value of the slot. For more information, see [PutSlotType](#).

Type: String

Valid Values: ORIGINAL_VALUE | TOP_RESOLUTION

version

The version of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: \ \$LATEST | [0-9]+

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetSlotTypes

Service: Amazon Lex Model Building Service

Returns slot type information as follows:

- If you specify the `nameContains` field, returns the \$LATEST version of all slot types that contain the specified string.
- If you don't specify the `nameContains` field, returns information about the \$LATEST version of all slot types.

The operation requires permission for the `lex:GetSlotTypes` action.

Request Syntax

```
GET /slottypes/?maxResults=maxResults&nameContains=nameContains&nextToken=nextToken
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of slot types to return in the response. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

nameContains

Substring to match in slot type names. A slot type will be returned if any part of its name matches the substring. For example, "xyz" matches both "xyzabc" and "abcxyz."

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

nextToken

A pagination token that fetches the next page of slot types. If the response to this API call is truncated, Amazon Lex returns a pagination token in the response. To fetch next page of slot types, specify the pagination token in the next request.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "slotTypes": [
    {
      "createdDate": number,
      "description": "string",
      "lastUpdatedDate": number,
      "name": "string",
      "version": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[nextToken](#)

If the response is truncated, it includes a pagination token that you can specify in your next request to fetch the next page of slot types.

Type: String

[slotTypes](#)

An array of objects, one for each slot type, that provides information such as the name of the slot type, the version, and a description.

Type: Array of [SlotTypeMetadata](#) objects

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

GetSlotTypeVersions

Service: Amazon Lex Model Building Service

Gets information about all versions of a slot type.

The GetSlotTypeVersions operation returns a SlotTypeMetadata object for each version of a slot type. For example, if a slot type has three numbered versions, the GetSlotTypeVersions operation returns four SlotTypeMetadata objects in the response, one for each numbered version and one for the \$LATEST version.

The GetSlotTypeVersions operation always returns at least one version, the \$LATEST version.

This operation requires permissions for the `lex:GetSlotTypeVersions` action.

Request Syntax

```
GET /slottypes/name/versions/?maxResults=maxResults&nextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of slot type versions to return in the response. The default is 10.

Valid Range: Minimum value of 1. Maximum value of 50.

name

The name of the slot type for which versions should be returned.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

nextToken

A pagination token for fetching the next page of slot type versions. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of versions, specify the pagination token in the next request.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "slotTypes": [
    {
      "createdDate": number,
      "description": "string",
      "lastUpdatedDate": number,
      "name": "string",
      "version": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

nextToken

A pagination token for fetching the next page of slot type versions. If the response to this call is truncated, Amazon Lex returns a pagination token in the response. To fetch the next page of versions, specify the pagination token in the next request.

Type: String

slotTypes

An array of `SlotTypeMetadata` objects, one for each numbered version of the slot type plus one for the `$LATEST` version.

Type: Array of [SlotTypeMetadata](#) objects

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

GetUtterancesView

Service: Amazon Lex Model Building Service

Use the `GetUtterancesView` operation to get information about the utterances that your users have made to your bot. You can use this list to tune the utterances that your bot responds to.

For example, say that you have created a bot to order flowers. After your users have used your bot for a while, use the `GetUtterancesView` operation to see the requests that they have made and whether they have been successful. You might find that the utterance "I want flowers" is not being recognized. You could add this utterance to the `OrderFlowers` intent so that your bot recognizes that utterance.

After you publish a new version of a bot, you can get information about the old version and the new so that you can compare the performance across the two versions.

Utterance statistics are generated once a day. Data is available for the last 15 days. You can request information for up to 5 versions of your bot in each request. Amazon Lex returns the most frequent utterances received by the bot in the last 15 days. The response contains information about a maximum of 100 utterances for each version.

Utterance statistics are not generated under the following conditions:

- The `childDirected` field was set to true when the bot was created.
- You are using slot obfuscation with one or more slots.
- You opted out of participating in improving Amazon Lex.

This operation requires permissions for the `lex:GetUtterancesView` action.

Request Syntax

```
GET /bots/botname/utterances?
view=aggregation&bot_versions=botVersions&status_type=statusType HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

botname

The name of the bot for which utterance information should be returned.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z]_?)+\$

Required: Yes

botVersions

An array of bot versions for which utterance information should be returned. The limit is 5 versions per request.

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: \\$(LATEST|[0-9])+

Required: Yes

statusType

To return utterances that were recognized and handled, use Detected. To return utterances that were not recognized, use Missed.

Valid Values: Detected | Missed

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "botName": "string",
  "utterances": [
    {
      "botVersion": "string",
      "utterances": [
```

```
{
  {
    {
      "count": number,
      "distinctUsers": number,
      "firstUtteredDate": number,
      "lastUtteredDate": number,
      "utteranceString": "string"
    }
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

botName

The name of the bot for which utterance information was returned.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z]_?)+\$

utterances

An array of [UtteranceList](#) objects, each containing a list of [UtteranceData](#) objects describing the utterances that were processed by your bot. The response contains a maximum of 100 [UtteranceData](#) objects for each version. Amazon Lex returns the most frequent utterances received by the bot in the last 15 days.

Type: Array of [UtteranceList](#) objects

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Service: Amazon Lex Model Building Service

Gets a list of tags associated with the specified resource. Only bots, bot aliases, and bot channels can have tags associated with them.

Request Syntax

```
GET /tags/resourceArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the resource to get a list of tags for.

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

tags

The tags associated with a resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PutBot

Service: Amazon Lex Model Building Service

Creates an Amazon Lex conversational bot or replaces an existing bot. When you create or update a bot you are only required to specify a name, a locale, and whether the bot is directed toward children under age 13. You can use this to add intents later, or to remove intents from an existing bot. When you create a bot with the minimum information, the bot is created or updated but Amazon Lex returns the response `FAILED`. You can build the bot after you add one or more intents. For more information about Amazon Lex bots, see [Amazon Lex: How It Works](#).

If you specify the name of an existing bot, the fields in the request replace the existing values in the `$LATEST` version of the bot. Amazon Lex removes any fields that you don't provide values for in the request, except for the `idleTTLInSeconds` and `privacySettings` fields, which are set to their default values. If you don't specify values for required fields, Amazon Lex throws an exception.

This operation requires permissions for the `lex:PutBot` action. For more information, see [Identity and Access Management for Amazon Lex](#).

Request Syntax

```
PUT /bots/name/versions/$LATEST HTTP/1.1
Content-type: application/json
```

```
{
  "abortStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupName": number
      }
    ],
    "responseCard": "string"
  },
  "checksum": "string",
  "childDirected": boolean,
  "clarificationPrompt": {
    "maxAttempts": number,
    "messages": [
      {
```

```

        "content": "string",
        "contentType": "string",
        "groupNumber": number
    }
],
    "responseCard": "string"
},
    "createVersion": boolean,
    "description": "string",
    "detectSentiment": boolean,
    "enableModelImprovements": boolean,
    "idleSessionTTLInSeconds": number,
    "intents": [
        {
            "intentName": "string",
            "intentVersion": "string"
        }
    ],
    "locale": "string",
    "nluIntentConfidenceThreshold": number,
    "processBehavior": "string",
    "tags": [
        {
            "key": "string",
            "value": "string"
        }
    ],
    "voiceId": "string"
}

```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the bot. The name is *not* case sensitive.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

Request Body

The request accepts the following data in JSON format.

abortStatement

When Amazon Lex can't understand the user's input in context, it tries to elicit the information a few times. After that, Amazon Lex sends the message defined in `abortStatement` to the user, and then cancels the conversation. To set the number of retries, use the `valueElicitationPrompt` field for the slot type.

For example, in a pizza ordering bot, Amazon Lex might ask a user "What type of crust would you like?" If the user's response is not one of the expected responses (for example, "thin crust, "deep dish," etc.), Amazon Lex tries to elicit a correct response a few more times.

For example, in a pizza ordering application, `OrderPizza` might be one of the intents. This intent might require the `CrustType` slot. You specify the `valueElicitationPrompt` field when you create the `CrustType` slot.

If you have defined a fallback intent the cancel statement will not be sent to the user, the fallback intent is used instead. For more information, see [AMAZON.FallbackIntent](#).

Type: [Statement](#) object

Required: No

checksum

Identifies a specific revision of the \$LATEST version.

When you create a new bot, leave the checksum field blank. If you specify a checksum you get a `BadRequestException` exception.

When you want to update a bot, set the checksum field to the checksum of the most recent revision of the \$LATEST version. If you don't specify the checksum field, or if the checksum does not match the \$LATEST version, you get a `PreconditionFailedException` exception.

Type: String

Required: No

childDirected

For each Amazon Lex bot created with the Amazon Lex Model Building Service, you must specify whether your use of Amazon Lex is related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to the Children's Online Privacy Protection Act (COPPA) by specifying `true` or `false` in the `childDirected` field. By specifying `true` in the `childDirected` field, you confirm that your use of Amazon Lex **is** related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA. By specifying `false` in the `childDirected` field, you confirm that your use of Amazon Lex **is not** related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA. You may not specify a default value for the `childDirected` field that does not accurately reflect whether your use of Amazon Lex is related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA.

If your use of Amazon Lex relates to a website, program, or other application that is directed in whole or in part, to children under age 13, you must obtain any required verifiable parental consent under COPPA. For information regarding the use of Amazon Lex in connection with websites, programs, or other applications that are directed or targeted, in whole or in part, to children under age 13, see the [Amazon Lex FAQ](#).

Type: Boolean

Required: Yes

clarificationPrompt

When Amazon Lex doesn't understand the user's intent, it uses this message to get clarification. To specify how many times Amazon Lex should repeat the clarification prompt, use the `maxAttempts` field. If Amazon Lex still doesn't understand, it sends the message in the `abortStatement` field.

When you create a clarification prompt, make sure that it suggests the correct response from the user. For example, for a bot that orders pizza and drinks, you might create this clarification prompt: "What would you like to do? You can say 'Order a pizza' or 'Order a drink.'"

If you have defined a fallback intent, it will be invoked if the clarification prompt is repeated the number of times defined in the `maxAttempts` field. For more information, see [AMAZON.FallbackIntent](#).

If you don't define a clarification prompt, at runtime Amazon Lex will return a 400 Bad Request exception in three cases:

- Follow-up prompt - When the user responds to a follow-up prompt but does not provide an intent. For example, in response to a follow-up prompt that says "Would you like anything else today?" the user says "Yes." Amazon Lex will return a 400 Bad Request exception because it does not have a clarification prompt to send to the user to get an intent.
- Lambda function - When using a Lambda function, you return an `ElicitIntent` dialog type. Since Amazon Lex does not have a clarification prompt to get an intent from the user, it returns a 400 Bad Request exception.
- PutSession operation - When using the PutSession operation, you send an `ElicitIntent` dialog type. Since Amazon Lex does not have a clarification prompt to get an intent from the user, it returns a 400 Bad Request exception.

Type: [Prompt](#) object

Required: No

[**createVersion**](#)

When set to true a new numbered version of the bot is created. This is the same as calling the `CreateBotVersion` operation. If you don't specify `createVersion`, the default is false.

Type: Boolean

Required: No

[**description**](#)

A description of the bot.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

[**detectSentiment**](#)

When set to true user utterances are sent to Amazon Comprehend for sentiment analysis. If you don't specify `detectSentiment`, the default is false.

Type: Boolean

Required: No

[enableModelImprovements](#)

Set to `true` to enable access to natural language understanding improvements.

When you set the `enableModelImprovements` parameter to `true` you can use the `nluIntentConfidenceThreshold` parameter to configure confidence scores. For more information, see [Confidence Scores](#).

You can only set the `enableModelImprovements` parameter in certain Regions. If you set the parameter to `true`, your bot has access to accuracy improvements.

The Regions where you can set the `enableModelImprovements` parameter to `false` for the en-US locale are:

- US East (N. Virginia) (us-east-1)
- US West (Oregon) (us-west-2)
- Asia Pacific (Sydney) (ap-southeast-2)
- EU (Ireland) (eu-west-1)

In other Regions and locales, the `enableModelImprovements` parameter is set to `true` by default. In these Regions and locales setting the parameter to `false` throws a `ValidationException` exception.

Type: Boolean

Required: No

[idleSessionTTLInSeconds](#)

The maximum time in seconds that Amazon Lex retains the data gathered in a conversation.

A user interaction session remains active for the amount of time specified. If no conversation occurs during this time, the session expires and Amazon Lex deletes any data provided before the timeout.

For example, suppose that a user chooses the `OrderPizza` intent, but gets sidetracked halfway through placing an order. If the user doesn't complete the order within the specified time, Amazon Lex discards the slot information that it gathered, and the user must start over.

If you don't include the `idleSessionTTLInSeconds` element in a `PutBot` operation request, Amazon Lex uses the default value. This is also true if the request replaces an existing bot.

The default is 300 seconds (5 minutes).

Type: Integer

Valid Range: Minimum value of 60. Maximum value of 86400.

Required: No

intents

An array of Intent objects. Each intent represents a command that a user can express. For example, a pizza ordering bot might support an OrderPizza intent. For more information, see [Amazon Lex: How It Works](#).

Type: Array of [Intent](#) objects

Required: No

locale

Specifies the target locale for the bot. Any intent used in the bot must be compatible with the locale of the bot.

The default is en-US.

Type: String

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

Required: Yes

nluIntentConfidenceThreshold

Determines the threshold where Amazon Lex will insert the AMAZON.FallbackIntent, AMAZON.KendraSearchIntent, or both when returning alternative intents in a [PostContent](#) or [PostText](#) response. AMAZON.FallbackIntent and AMAZON.KendraSearchIntent are only inserted if they are configured for the bot.

You must set the enableModelImprovements parameter to true to use confidence scores in the following regions.

- US East (N. Virginia) (us-east-1)
- US West (Oregon) (us-west-2)

- Asia Pacific (Sydney) (ap-southeast-2)
- EU (Ireland) (eu-west-1)

In other Regions, the `enableModelImprovements` parameter is set to `true` by default.

For example, suppose a bot is configured with the confidence threshold of 0.80 and the `AMAZON.FallbackIntent`. Amazon Lex returns three alternative intents with the following confidence scores: `IntentA` (0.70), `IntentB` (0.60), `IntentC` (0.50). The response from the `PostText` operation would be:

- `AMAZON.FallbackIntent`
- `IntentA`
- `IntentB`
- `IntentC`

Type: Double

Valid Range: Minimum value of 0. Maximum value of 1.

Required: No

[processBehavior](#)

If you set the `processBehavior` element to `BUILD`, Amazon Lex builds the bot so that it can be run. If you set the element to `SAVE` Amazon Lex saves the bot, but doesn't build it.

If you don't specify this value, the default value is `BUILD`.

Type: String

Valid Values: `SAVE` | `BUILD`

Required: No

[tags](#)

A list of tags to add to the bot. You can only add tags when you create a bot, you can't use the `PutBot` operation to update the tags on a bot. To update tags, use the `TagResource` operation.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: No

voiceId

The Amazon Polly voice ID that you want Amazon Lex to use for voice interactions with the user. The locale configured for the voice must match the locale of the bot. For more information, see [Voices in Amazon Polly](#) in the *Amazon Polly Developer Guide*.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "abortStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "checksum": "string",
  "childDirected": boolean,
  "clarificationPrompt": {
    "maxAttempts": number,
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
}
```

```

"createdDate": number,
"createVersion": boolean,
"description": "string",
"detectSentiment": boolean,
"enableModelImprovements": boolean,
"failureReason": "string",
"idleSessionTTLInSeconds": number,
"intents": [
  {
    "intentName": "string",
    "intentVersion": "string"
  }
],
"lastUpdatedDate": number,
"locale": "string",
"name": "string",
"nluIntentConfidenceThreshold": number,
"status": "string",
"tags": [
  {
    "key": "string",
    "value": "string"
  }
],
"version": "string",
"voiceId": "string"
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

abortStatement

The message that Amazon Lex uses to cancel a conversation. For more information, see [PutBot](#).

Type: [Statement](#) object

checksum

Checksum of the bot that you created.

Type: String

childDirected

For each Amazon Lex bot created with the Amazon Lex Model Building Service, you must specify whether your use of Amazon Lex is related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to the Children's Online Privacy Protection Act (COPPA) by specifying `true` or `false` in the `childDirected` field. By specifying `true` in the `childDirected` field, you confirm that your use of Amazon Lex **is** related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA. By specifying `false` in the `childDirected` field, you confirm that your use of Amazon Lex **is not** related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA. You may not specify a default value for the `childDirected` field that does not accurately reflect whether your use of Amazon Lex is related to a website, program, or other application that is directed or targeted, in whole or in part, to children under age 13 and subject to COPPA.

If your use of Amazon Lex relates to a website, program, or other application that is directed in whole or in part, to children under age 13, you must obtain any required verifiable parental consent under COPPA. For information regarding the use of Amazon Lex in connection with websites, programs, or other applications that are directed or targeted, in whole or in part, to children under age 13, see the [Amazon Lex FAQ](#).

Type: Boolean

clarificationPrompt

The prompts that Amazon Lex uses when it doesn't understand the user's intent. For more information, see [PutBot](#).

Type: [Prompt](#) object

createdDate

The date that the bot was created.

Type: Timestamp

createVersion

`True` if a new version of the bot was created. If the `createVersion` field was not specified in the request, the `createVersion` field is set to `false` in the response.

Type: Boolean

description

A description of the bot.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

detectSentiment

true if the bot is configured to send user utterances to Amazon Comprehend for sentiment analysis. If the detectSentiment field was not specified in the request, the detectSentiment field is false in the response.

Type: Boolean

enableModelImprovements

Indicates whether the bot uses accuracy improvements. true indicates that the bot is using the improvements, otherwise, false.

Type: Boolean

failureReason

If status is FAILED, Amazon Lex provides the reason that it failed to build the bot.

Type: String

idleSessionTTLInSeconds

The maximum length of time that Amazon Lex retains the data gathered in a conversation. For more information, see [PutBot](#).

Type: Integer

Valid Range: Minimum value of 60. Maximum value of 86400.

intents

An array of Intent objects. For more information, see [PutBot](#).

Type: Array of [Intent](#) objects

lastUpdatedDate

The date that the bot was updated. When you create a resource, the creation date and last updated date are the same.

Type: Timestamp

locale

The target locale for the bot.

Type: String

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

name

The name of the bot.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z]_?)+\$

nlIntentConfidenceThreshold

The score that determines where Amazon Lex inserts the `AMAZON.FallbackIntent`, `AMAZON.KendraSearchIntent`, or both when returning alternative intents in a [PostContent](#) or [PostText](#) response. `AMAZON.FallbackIntent` is inserted if the confidence score for all intents is below this value. `AMAZON.KendraSearchIntent` is only inserted if it is configured for the bot.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 1.

status

When you send a request to create a bot with `processBehavior` set to `BUILD`, Amazon Lex sets the `status` response element to `BUILDING`.

In the `READY_BASIC_TESTING` state you can test the bot with user inputs that exactly match the utterances configured for the bot's intents and values in the slot types.

If Amazon Lex can't build the bot, Amazon Lex sets `status` to `FAILED`. Amazon Lex returns the reason for the failure in the `failureReason` response element.

When you set `processBehavior` to `SAVE`, Amazon Lex sets the status code to `NOT_BUILT`.

When the bot is in the `READY` state you can test and publish the bot.

Type: String

Valid Values: `BUILDING` | `READY` | `READY_BASIC_TESTING` | `FAILED` | `NOT_BUILT`

tags

A list of tags associated with the bot.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

version

The version of the bot. For a new bot, the version is always `$LATEST`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

voiceId

The Amazon Polly voice ID that Amazon Lex uses for voice interaction with the user. For more information, see [PutBot](#).

Type: String

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

PreconditionFailedException

The checksum of the resource that you are trying to change does not match the checksum in the request. Check the resource's checksum and try again.

HTTP Status Code: 412

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PutBotAlias

Service: Amazon Lex Model Building Service

Creates an alias for the specified version of the bot or replaces an alias for the specified bot. To change the version of the bot that the alias points to, replace the alias. For more information about aliases, see [Versioning and Aliases](#).

This operation requires permissions for the `lex:PutBotAlias` action.

Request Syntax

```
PUT /bots/botName/aliases/name HTTP/1.1
Content-type: application/json
```

```
{
  "botVersion": "string",
  "checksum": "string",
  "conversationLogs": {
    "iamRoleArn": "string",
    "logSettings": [
      {
        "destination": "string",
        "kmsKeyArn": "string",
        "logType": "string",
        "resourceArn": "string"
      }
    ]
  },
  "description": "string",
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

URI Request Parameters

The request uses the following URI parameters.

botName

The name of the bot.

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

name

The name of the alias. The name is *not* case sensitive.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request accepts the following data in JSON format.

botVersion

The version of the bot.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: Yes

checksum

Identifies a specific revision of the \$LATEST version.

When you create a new bot alias, leave the checksum field blank. If you specify a checksum you get a `BadRequestException` exception.

When you want to update a bot alias, set the checksum field to the checksum of the most recent revision of the \$LATEST version. If you don't specify the checksum field, or if the checksum does not match the \$LATEST version, you get a `PreconditionFailedException` exception.

Type: String

Required: No

conversationLogs

Settings for conversation logs for the alias.

Type: [ConversationLogsRequest](#) object

Required: No

description

A description of the alias.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

tags

A list of tags to add to the bot alias. You can only add tags when you create an alias, you can't use the `PutBotAlias` operation to update the tags on a bot alias. To update tags, use the `TagResource` operation.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "botName": "string",
  "botVersion": "string",
  "checksum": "string",
  "conversationLogs": {
    "iamRoleArn": "string",
    "logSettings": [
      {
        "destination": "string",
        "kmsKeyArn": "string",
        "logType": "string",
        "resourceArn": "string",
        "resourcePrefix": "string"
      }
    ]
  },
  "createdDate": number,
  "description": "string",
  "lastUpdatedDate": number,
  "name": "string",
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

botName

The name of the bot that the alias points to.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

botVersion

The version of the bot that the alias points to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

checksum

The checksum for the current version of the alias.

Type: String

conversationLogs

The settings that determine how Amazon Lex uses conversation logs for the alias.

Type: [ConversationLogsResponse](#) object

createdDate

The date that the bot alias was created.

Type: Timestamp

description

A description of the alias.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

lastUpdatedDate

The date that the bot alias was updated. When you create a resource, the creation date and the last updated date are the same.

Type: Timestamp

name

The name of the alias.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

tags

A list of tags associated with a bot.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

PreconditionFailedException

The checksum of the resource that you are trying to change does not match the checksum in the request. Check the resource's checksum and try again.

HTTP Status Code: 412

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PutIntent

Service: Amazon Lex Model Building Service

Creates an intent or replaces an existing intent.

To define the interaction between the user and your bot, you use one or more intents. For a pizza ordering bot, for example, you would create an `OrderPizza` intent.

To create an intent or replace an existing intent, you must provide the following:

- Intent name. For example, `OrderPizza`.
- Sample utterances. For example, "Can I order a pizza, please." and "I want to order a pizza."
- Information to be gathered. You specify slot types for the information that your bot will request from the user. You can specify standard slot types, such as a date or a time, or custom slot types such as the size and crust of a pizza.
- How the intent will be fulfilled. You can provide a Lambda function or configure the intent to return the intent information to the client application. If you use a Lambda function, when all of the intent information is available, Amazon Lex invokes your Lambda function. If you configure your intent to return the intent information to the client application.

You can specify other optional information in the request, such as:

- A confirmation prompt to ask the user to confirm an intent. For example, "Shall I order your pizza?"
- A conclusion statement to send to the user after the intent has been fulfilled. For example, "I placed your pizza order."
- A follow-up prompt that asks the user for additional activity. For example, asking "Do you want to order a drink with your pizza?"

If you specify an existing intent name to update the intent, Amazon Lex replaces the values in the `$LATEST` version of the intent with the values in the request. Amazon Lex removes fields that you don't provide in the request. If you don't specify the required fields, Amazon Lex throws an exception. When you update the `$LATEST` version of an intent, the status field of any bot that uses the `$LATEST` version of the intent is set to `NOT_BUILT`.

For more information, see [Amazon Lex: How It Works](#).

This operation requires permissions for the `lex:PutIntent` action.

Request Syntax

```
PUT /intents/name/versions/$LATEST HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "checksum": "string",
  "conclusionStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "confirmationPrompt": {
    "maxAttempts": number,
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "createVersion": boolean,
  "description": "string",
  "dialogCodeHook": {
    "messageVersion": "string",
    "uri": "string"
  },
  "followUpPrompt": {
    "prompt": {
      "maxAttempts": number,
      "messages": [
        {
          "content": "string",
          "contentType": "string",
          "groupNumber": number
        }
      ]
    },
  ],
}
```

```
    "responseCard": "string"
  },
  "rejectionStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  }
},
"fulfillmentActivity": {
  "codeHook": {
    "messageVersion": "string",
    "uri": "string"
  },
  "type": "string"
},
inputContexts": [
  {
    "name": "string"
  }
],
"kendraConfiguration": {
  "kendraIndex": "string",
  "queryFilterString": "string",
  "role": "string"
},
outputContexts": [
  {
    "name": "string",
    "timeToLiveInSeconds": number,
    "turnsToLive": number
  }
],
"parentIntentSignature": "string",
"rejectionStatement": {
  "messages": [
    {
      "content": "string",
      "contentType": "string",
      "groupNumber": number
    }
  ]
}
```

```

    }
  ],
  "responseCard": "string"
},
"sampleUtterances": [ "string" ],
"slots": [
  {
    "defaultValueSpec": {
      "defaultValueList": [
        {
          "defaultValue": "string"
        }
      ]
    },
    "description": "string",
    "name": "string",
    "obfuscationSetting": "string",
    "priority": number,
    "responseCard": "string",
    "sampleUtterances": [ "string" ],
    "slotConstraint": "string",
    "slotType": "string",
    "slotTypeVersion": "string",
    "valueElicitationPrompt": {
      "maxAttempts": number,
      "messages": [
        {
          "content": "string",
          "contentType": "string",
          "groupNumber": number
        }
      ]
    },
    "responseCard": "string"
  }
]
}

```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the intent. The name is *not* case sensitive.

The name can't match a built-in intent name, or a built-in intent name with "AMAZON." removed. For example, because there is a built-in intent called `AMAZON.HelpIntent`, you can't create a custom intent called `HelpIntent`.

For a list of built-in intents, see [Standard Built-in Intents](#) in the *Alexa Skills Kit*.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

Request Body

The request accepts the following data in JSON format.

checksum

Identifies a specific revision of the `$LATEST` version.

When you create a new intent, leave the checksum field blank. If you specify a checksum you get a `BadRequestException` exception.

When you want to update a intent, set the checksum field to the checksum of the most recent revision of the `$LATEST` version. If you don't specify the checksum field, or if the checksum does not match the `$LATEST` version, you get a `PreconditionFailedException` exception.

Type: String

Required: No

conclusionStatement

The statement that you want Amazon Lex to convey to the user after the intent is successfully fulfilled by the Lambda function.

This element is relevant only if you provide a Lambda function in the `fulfillmentActivity`. If you return the intent to the client application, you can't specify this element.

Note

The `followUpPrompt` and `conclusionStatement` are mutually exclusive. You can specify only one.

Type: [Statement](#) object

Required: No

[confirmationPrompt](#)

Prompts the user to confirm the intent. This question should have a yes or no answer.

Amazon Lex uses this prompt to ensure that the user acknowledges that the intent is ready for fulfillment. For example, with the `OrderPizza` intent, you might want to confirm that the order is correct before placing it. For other intents, such as intents that simply respond to user questions, you might not need to ask the user for confirmation before providing the information.

Note

You must provide both the `rejectionStatement` and the `confirmationPrompt`, or neither.

Type: [Prompt](#) object

Required: No

[createVersion](#)

When set to `true` a new numbered version of the intent is created. This is the same as calling the `CreateIntentVersion` operation. If you do not specify `createVersion`, the default is `false`.

Type: Boolean

Required: No

[description](#)

A description of the intent.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

[dialogCodeHook](#)

Specifies a Lambda function to invoke for each user input. You can invoke this Lambda function to personalize user interaction.

For example, suppose your bot determines that the user is John. Your Lambda function might retrieve John's information from a backend database and prepopulate some of the values. For example, if you find that John is gluten intolerant, you might set the corresponding intent slot, `GlutenIntolerant`, to true. You might find John's phone number and set the corresponding session attribute.

Type: [CodeHook](#) object

Required: No

[followUpPrompt](#)

Amazon Lex uses this prompt to solicit additional activity after fulfilling an intent. For example, after the `OrderPizza` intent is fulfilled, you might prompt the user to order a drink.

The action that Amazon Lex takes depends on the user's response, as follows:

- If the user says "Yes" it responds with the clarification prompt that is configured for the bot.
- if the user says "Yes" and continues with an utterance that triggers an intent it starts a conversation for the intent.
- If the user says "No" it responds with the rejection statement configured for the the follow-up prompt.
- If it doesn't recognize the utterance it repeats the follow-up prompt again.

The `followUpPrompt` field and the `conclusionStatement` field are mutually exclusive. You can specify only one.

Type: [FollowUpPrompt](#) object

Required: No

fulfillmentActivity

Required. Describes how the intent is fulfilled. For example, after a user provides all of the information for a pizza order, `fulfillmentActivity` defines how the bot places an order with a local pizza store.

You might configure Amazon Lex to return all of the intent information to the client application, or direct it to invoke a Lambda function that can process the intent (for example, place an order with a pizzeria).

Type: [FulfillmentActivity](#) object

Required: No

inputContexts

An array of `InputContext` objects that lists the contexts that must be active for Amazon Lex to choose the intent in a conversation with the user.

Type: Array of [InputContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 5 items.

Required: No

kendraConfiguration

Configuration information required to use the `AMAZON.KendraSearchIntent` intent to connect to an Amazon Kendra index. For more information, see [AMAZON.KendraSearchIntent](#).

Type: [KendraConfiguration](#) object

Required: No

outputContexts

An array of `OutputContext` objects that lists the contexts that the intent activates when the intent is fulfilled.

Type: Array of [OutputContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

parentIntentSignature

A unique identifier for the built-in intent to base this intent on. To find the signature for an intent, see [Standard Built-in Intents](#) in the *Alexa Skills Kit*.

Type: String

Required: No

rejectionStatement

When the user answers "no" to the question defined in `confirmationPrompt`, Amazon Lex responds with this statement to acknowledge that the intent was canceled.

Note

You must provide both the `rejectionStatement` and the `confirmationPrompt`, or neither.

Type: [Statement](#) object

Required: No

sampleUtterances

An array of utterances (strings) that a user might say to signal the intent. For example, "I want {PizzaSize} pizza", "Order {Quantity} {PizzaSize} pizzas".

In each utterance, a slot name is enclosed in curly braces.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 1500 items.

Length Constraints: Minimum length of 1. Maximum length of 200.

Required: No

slots

An array of intent slots. At runtime, Amazon Lex elicits required slot values from the user using prompts defined in the slots. For more information, see [Amazon Lex: How It Works](#).

Type: Array of [Slot](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "checksum": "string",
  "conclusionStatement": {
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "confirmationPrompt": {
    "maxAttempts": number,
    "messages": [
      {
        "content": "string",
        "contentType": "string",
        "groupNumber": number
      }
    ],
    "responseCard": "string"
  },
  "createdDate": number,
  "createVersion": boolean,
  "description": "string",
  "dialogCodeHook": {
    "messageVersion": "string",
    "uri": "string"
  },
  "followUpPrompt": {
    "prompt": {
      "maxAttempts": number,
      "messages": [
```

```

        {
            "content": "string",
            "contentType": "string",
            "groupNumber": number
        }
    ],
    "responseCard": "string"
},
"rejectionStatement": {
    "messages": [
        {
            "content": "string",
            "contentType": "string",
            "groupNumber": number
        }
    ],
    "responseCard": "string"
}
},
"fulfillmentActivity": {
    "codeHook": {
        "messageVersion": "string",
        "uri": "string"
    },
    "type": "string"
},
"inputContexts": [
    {
        "name": "string"
    }
],
"kendraConfiguration": {
    "kendraIndex": "string",
    "queryFilterString": "string",
    "role": "string"
},
"lastUpdatedDate": number,
"name": "string",
"outputContexts": [
    {
        "name": "string",
        "timeToLiveInSeconds": number,
        "turnsToLive": number
    }
]
}

```

```

],
"parentIntentSignature": "string",
"rejectionStatement": {
  "messages": [
    {
      "content": "string",
      "contentType": "string",
      "groupNumber": number
    }
  ],
  "responseCard": "string"
},
"sampleUtterances": [ "string" ],
"slots": [
  {
    "defaultValueSpec": {
      "defaultValueList": [
        {
          "defaultValue": "string"
        }
      ]
    },
    "description": "string",
    "name": "string",
    "obfuscationSetting": "string",
    "priority": number,
    "responseCard": "string",
    "sampleUtterances": [ "string" ],
    "slotConstraint": "string",
    "slotType": "string",
    "slotTypeVersion": "string",
    "valueElicitationPrompt": {
      "maxAttempts": number,
      "messages": [
        {
          "content": "string",
          "contentType": "string",
          "groupNumber": number
        }
      ]
    },
    "responseCard": "string"
  }
]
],

```

```
"version": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[checksum](#)

Checksum of the `$LATEST` version of the intent created or updated.

Type: String

[conclusionStatement](#)

After the Lambda function specified in the `fulfillmentActivity` intent fulfills the intent, Amazon Lex conveys this statement to the user.

Type: [Statement](#) object

[confirmationPrompt](#)

If defined in the intent, Amazon Lex prompts the user to confirm the intent before fulfilling it.

Type: [Prompt](#) object

[createdDate](#)

The date that the intent was created.

Type: Timestamp

[createVersion](#)

True if a new version of the intent was created. If the `createVersion` field was not specified in the request, the `createVersion` field is set to false in the response.

Type: Boolean

[description](#)

A description of the intent.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

dialogCodeHook

If defined in the intent, Amazon Lex invokes this Lambda function for each user input.

Type: [CodeHook](#) object

followUpPrompt

If defined in the intent, Amazon Lex uses this prompt to solicit additional user activity after the intent is fulfilled.

Type: [FollowUpPrompt](#) object

fulfillmentActivity

If defined in the intent, Amazon Lex invokes this Lambda function to fulfill the intent after the user provides all of the information required by the intent.

Type: [FulfillmentActivity](#) object

inputContexts

An array of InputContext objects that lists the contexts that must be active for Amazon Lex to choose the intent in a conversation with the user.

Type: Array of [InputContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 5 items.

kendraConfiguration

Configuration information, if any, required to connect to an Amazon Kendra index and use the AMAZON.KendraSearchIntent intent.

Type: [KendraConfiguration](#) object

lastUpdatedDate

The date that the intent was updated. When you create a resource, the creation date and last update dates are the same.

Type: Timestamp

name

The name of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

outputContexts

An array of `OutputContext` objects that lists the contexts that the intent activates when the intent is fulfilled.

Type: Array of [OutputContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

parentIntentSignature

A unique identifier for the built-in intent that this intent is based on.

Type: String

rejectionStatement

If the user answers "no" to the question defined in `confirmationPrompt` Amazon Lex responds with this statement to acknowledge that the intent was canceled.

Type: [Statement](#) object

sampleUtterances

An array of sample utterances that are configured for the intent.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 1500 items.

Length Constraints: Minimum length of 1. Maximum length of 200.

slots

An array of intent slots that are configured for the intent.

Type: Array of [Slot](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

version

The version of the intent. For a new intent, the version is always \$LATEST.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: \ \$LATEST | [0-9]+

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

PreconditionFailedException

The checksum of the resource that you are trying to change does not match the checksum in the request. Check the resource's checksum and try again.

HTTP Status Code: 412

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PutSlotType

Service: Amazon Lex Model Building Service

Creates a custom slot type or replaces an existing custom slot type.

To create a custom slot type, specify a name for the slot type and a set of enumeration values, which are the values that a slot of this type can assume. For more information, see [Amazon Lex: How It Works](#).

If you specify the name of an existing slot type, the fields in the request replace the existing values in the \$LATEST version of the slot type. Amazon Lex removes the fields that you don't provide in the request. If you don't specify required fields, Amazon Lex throws an exception. When you update the \$LATEST version of a slot type, if a bot uses the \$LATEST version of an intent that contains the slot type, the bot's status field is set to NOT_BUILT.

This operation requires permissions for the `lex:PutSlotType` action.

Request Syntax

```
PUT /slottypes/name/versions/$LATEST HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "checksum": "string",
  "createVersion": boolean,
  "description": "string",
  "enumerationValues": [
    {
      "synonyms": [ "string" ],
      "value": "string"
    }
  ],
  "parentSlotTypeSignature": "string",
  "slotTypeConfigurations": [
    {
      "regexConfiguration": {
        "pattern": "string"
      }
    }
  ],
  "valueSelectionStrategy": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the slot type. The name is *not* case sensitive.

The name can't match a built-in slot type name, or a built-in slot type name with "AMAZON." removed. For example, because there is a built-in slot type called AMAZON.DATE, you can't create a custom slot type called DATE.

For a list of built-in slot types, see [Slot Type Reference](#) in the *Alexa Skills Kit*.

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

Required: Yes

Request Body

The request accepts the following data in JSON format.

[checksum](#)

Identifies a specific revision of the \$LATEST version.

When you create a new slot type, leave the checksum field blank. If you specify a checksum you get a `BadRequestException` exception.

When you want to update a slot type, set the checksum field to the checksum of the most recent revision of the \$LATEST version. If you don't specify the checksum field, or if the checksum does not match the \$LATEST version, you get a `PreconditionFailedException` exception.

Type: String

Required: No

createVersion

When set to true a new numbered version of the slot type is created. This is the same as calling the `CreateSlotTypeVersion` operation. If you do not specify `createVersion`, the default is false.

Type: Boolean

Required: No

description

A description of the slot type.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

enumerationValues

A list of `EnumerationValue` objects that defines the values that the slot type can take. Each value can have a list of synonyms, which are additional values that help train the machine learning model about the values that it resolves for a slot.

A regular expression slot type doesn't require enumeration values. All other slot types require a list of enumeration values.

When Amazon Lex resolves a slot value, it generates a resolution list that contains up to five possible values for the slot. If you are using a Lambda function, this resolution list is passed to the function. If you are not using a Lambda function you can choose to return the value that the user entered or the first value in the resolution list as the slot value. The `valueSelectionStrategy` field indicates the option to use.

Type: Array of [EnumerationValue](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10000 items.

Required: No

parentSlotTypeSignature

The built-in slot type used as the parent of the slot type. When you define a parent slot type, the new slot type has all of the same configuration as the parent.

Only `AMAZON.AlphaNumeric` is supported.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^((AMAZON\.)_?|[A-Za-z]_?)+`

Required: No

[slotTypeConfigurations](#)

Configuration information that extends the parent built-in slot type. The configuration is added to the settings for the parent slot type.

Type: Array of [SlotTypeConfiguration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

[valueSelectionStrategy](#)

Determines the slot resolution strategy that Amazon Lex uses to return slot type values. The field can be set to one of the following values:

- `ORIGINAL_VALUE` - Returns the value entered by the user, if the user value is similar to the slot value.
- `TOP_RESOLUTION` - If there is a resolution list for the slot, return the first value in the resolution list as the slot type value. If there is no resolution list, null is returned.

If you don't specify the `valueSelectionStrategy`, the default is `ORIGINAL_VALUE`.

Type: String

Valid Values: `ORIGINAL_VALUE` | `TOP_RESOLUTION`

Required: No

Response Syntax

```
HTTP/1.1 200
```

Content-type: application/json

```
{
  "checksum": "string",
  "createdDate": number,
  "createVersion": boolean,
  "description": "string",
  "enumerationValues": [
    {
      "synonyms": [ "string " ],
      "value": "string"
    }
  ],
  "lastUpdatedDate": number,
  "name": "string",
  "parentSlotTypeSignature": "string",
  "slotTypeConfigurations": [
    {
      "regexConfiguration": {
        "pattern": "string"
      }
    }
  ],
  "valueSelectionStrategy": "string",
  "version": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

checksum

Checksum of the \$LATEST version of the slot type.

Type: String

createdDate

The date that the slot type was created.

Type: Timestamp

createVersion

True if a new version of the slot type was created. If the createVersion field was not specified in the request, the createVersion field is set to false in the response.

Type: Boolean

description

A description of the slot type.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

enumerationValues

A list of EnumerationValue objects that defines the values that the slot type can take.

Type: Array of [EnumerationValue](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10000 items.

lastUpdatedDate

The date that the slot type was updated. When you create a slot type, the creation date and last update date are the same.

Type: Timestamp

name

The name of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: ^([A-Za-z]_?)+\$

parentSlotTypeSignature

The built-in slot type used as the parent of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^((AMAZON\. )_?|[A-Za-z]_?)+`

slotTypeConfigurations

Configuration information that extends the parent built-in slot type.

Type: Array of [SlotTypeConfiguration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

valueSelectionStrategy

The slot resolution strategy that Amazon Lex uses to determine the value of the slot. For more information, see [PutSlotType](#).

Type: String

Valid Values: ORIGINAL_VALUE | TOP_RESOLUTION

version

The version of the slot type. For a new slot type, the version is always \$LATEST.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

PreconditionFailedException

The checksum of the resource that you are trying to change does not match the checksum in the request. Check the resource's checksum and try again.

HTTP Status Code: 412

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

StartImport

Service: Amazon Lex Model Building Service

Starts a job to import a resource to Amazon Lex.

Request Syntax

```
POST /imports/ HTTP/1.1
Content-type: application/json

{
  "mergeStrategy": "string",
  "payload": blob,
  "resourceType": "string",
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

mergeStrategy

Specifies the action that the StartImport operation should take when there is an existing resource with the same name.

- **FAIL_ON_CONFLICT** - The import operation is stopped on the first conflict between a resource in the import file and an existing resource. The name of the resource causing the conflict is in the `failureReason` field of the response to the GetImport operation.

OVERWRITE_LATEST - The import operation proceeds even if there is a conflict with an existing resource. The **\$LATEST** version of the existing resource is overwritten with the data from the import file.

Type: String

Valid Values: OVERWRITE_LATEST | FAIL_ON_CONFLICT

Required: Yes

payload

A zip archive in binary format. The archive should contain one file, a JSON file containing the resource to import. The resource should match the type specified in the `resourceType` field.

Type: Base64-encoded binary data object

Required: Yes

resourceType

Specifies the type of resource to export. Each resource also exports any resources that it depends on.

- A bot exports dependent intents.
- An intent exports dependent slot types.

Type: String

Valid Values: BOT | INTENT | SLOT_TYPE

Required: Yes

tags

A list of tags to add to the imported bot. You can only add tags when you import a bot, you can't add tags to an intent or slot type.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: No

Response Syntax

```
HTTP/1.1 201
Content-type: application/json
```

```
{
  "createdDate": number,
  "importId": "string",
  "importStatus": "string",
  "mergeStrategy": "string",
  "name": "string",
  "resourceType": "string",
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 201 response.

The following data is returned in JSON format by the service.

createdDate

A timestamp for the date and time that the import job was requested.

Type: Timestamp

importId

The identifier for the specific import job.

Type: String

importStatus

The status of the import job. If the status is FAILED, you can get the reason for the failure using the `GetImport` operation.

Type: String

Valid Values: IN_PROGRESS | COMPLETE | FAILED

mergeStrategy

The action to take when there is a merge conflict.

Type: String

Valid Values: OVERWRITE_LATEST | FAIL_ON_CONFLICT

name

The name given to the import job.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: [a-zA-Z_]+

resourceType

The type of resource to import.

Type: String

Valid Values: BOT | INTENT | SLOT_TYPE

tags

A list of tags added to the imported bot.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

StartMigration

Service: Amazon Lex Model Building Service

Starts migrating a bot from Amazon Lex V1 to Amazon Lex V2. Migrate your bot when you want to take advantage of the new features of Amazon Lex V2.

For more information, see [Migrating a bot](#) in the *Amazon Lex developer guide*.

Request Syntax

```
POST /migrations HTTP/1.1
Content-type: application/json

{
  "migrationStrategy": "string",
  "v1BotName": "string",
  "v1BotVersion": "string",
  "v2BotName": "string",
  "v2BotRole": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[migrationStrategy](#)

The strategy used to conduct the migration.

- **CREATE_NEW** - Creates a new Amazon Lex V2 bot and migrates the Amazon Lex V1 bot to the new bot.
- **UPDATE_EXISTING** - Overwrites the existing Amazon Lex V2 bot metadata and the locale being migrated. It doesn't change any other locales in the Amazon Lex V2 bot. If the locale doesn't exist, a new locale is created in the Amazon Lex V2 bot.

Type: String

Valid Values: **CREATE_NEW** | **UPDATE_EXISTING**

Required: Yes

v1BotName

The name of the Amazon Lex V1 bot that you are migrating to Amazon Lex V2.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

v1BotVersion

The version of the bot to migrate to Amazon Lex V2. You can migrate the `$LATEST` version as well as any numbered version.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: Yes

v2BotName

The name of the Amazon Lex V2 bot that you are migrating the Amazon Lex V1 bot to.

- If the Amazon Lex V2 bot doesn't exist, you must use the `CREATE_NEW` migration strategy.
- If the Amazon Lex V2 bot exists, you must use the `UPDATE_EXISTING` migration strategy to change the contents of the Amazon Lex V2 bot.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[0-9a-zA-Z][_]?+$`

Required: Yes

v2BotRole

The IAM role that Amazon Lex uses to run the Amazon Lex V2 bot.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `^arn:[\w\ -]+:iam::[\d]{12}:role/.$`

Required: Yes

Response Syntax

```
HTTP/1.1 202
Content-type: application/json

{
  "migrationId": "string",
  "migrationStrategy": "string",
  "migrationTimestamp": number,
  "v1BotLocale": "string",
  "v1BotName": "string",
  "v1BotVersion": "string",
  "v2BotId": "string",
  "v2BotRole": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 202 response.

The following data is returned in JSON format by the service.

migrationId

The unique identifier that Amazon Lex assigned to the migration.

Type: String

Length Constraints: Fixed length of 10.

Pattern: `^[0-9a-zA-Z]+$`

migrationStrategy

The strategy used to conduct the migration.

Type: String

Valid Values: CREATE_NEW | UPDATE_EXISTING

migrationTimestamp

The date and time that the migration started.

Type: Timestamp

v1BotLocale

The locale used for the Amazon Lex V1 bot.

Type: String

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

v1BotName

The name of the Amazon Lex V1 bot that you are migrating to Amazon Lex V2.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z]_?)+\$

v1BotVersion

The version of the bot to migrate to Amazon Lex V2.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: \\$(LATEST|[0-9])+

v2BotId

The unique identifier for the Amazon Lex V2 bot.

Type: String

Length Constraints: Fixed length of 10.

Pattern: `^[0-9a-zA-Z]+$`

v2BotRole

The IAM role that Amazon Lex uses to run the Amazon Lex V2 bot.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `^arn:[\w\-.]+:iam::[\d]{12}:role/.+$`

Errors

AccessDeniedException

Your IAM user or role does not have permission to call the Amazon Lex V2 APIs required to migrate your bot.

HTTP Status Code: 403

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

TagResource

Service: Amazon Lex Model Building Service

Adds the specified tags to the specified resource. If a tag key already exists, the existing value is replaced with the new value.

Request Syntax

```
POST /tags/resourceArn HTTP/1.1
Content-type: application/json
```

```
{
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the bot, bot alias, or bot channel to tag.

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

Request Body

The request accepts the following data in JSON format.

tags

A list of tag keys to add to the resource. If a tag key already exists, the existing value is replaced with the new value.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: Yes

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UntagResource

Service: Amazon Lex Model Building Service

Removes tags from a bot, bot alias or bot channel.

Request Syntax

```
DELETE /tags/resourceArn?tagKeys=tagKeys HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the resource to remove the tags from.

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

tagKeys

A list of tag keys to remove from the resource. If a tag key does not exist on the resource, it is ignored.

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is not well formed. For example, a value is invalid or a required field is missing. Check the field values, and try again.

HTTP Status Code: 400

ConflictException

There was a conflict processing the request. Try your request again.

HTTP Status Code: 409

InternalFailureException

An internal Amazon Lex error occurred. Try your request again.

HTTP Status Code: 500

LimitExceededException

The request exceeded a limit. Try your request again.

HTTP Status Code: 429

NotFoundException

The resource specified in the request was not found. Check the resource and try again.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

Amazon Lex Runtime Service

The following actions are supported by Amazon Lex Runtime Service:

- [DeleteSession](#)
- [GetSession](#)
- [PostContent](#)
- [PostText](#)
- [PutSession](#)

DeleteSession

Service: Amazon Lex Runtime Service

Removes session information for a specified bot, alias, and user ID.

Request Syntax

```
DELETE /bot/botName/alias/botAlias/user/userId/session HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

botAlias

The alias in use for the bot that contains the session data.

Required: Yes

botName

The name of the bot that contains the session data.

Required: Yes

userId

The identifier of the user associated with the session data.

Length Constraints: Minimum length of 2. Maximum length of 100.

Pattern: [0-9a-zA-Z._:-]+

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
```

```
"botAlias": "string",  
"botName": "string",  
"sessionId": "string",  
"userId": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

botAlias

The alias in use for the bot associated with the session data.

Type: String

botName

The name of the bot associated with the session data.

Type: String

sessionId

The unique identifier for the session.

Type: String

userId

The ID of the client application user.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 100.

Pattern: [0-9a-zA-Z._:-]+

Errors

BadRequestException

Request validation failed, there is no usable message in the context, or the bot build failed, is still in progress, or contains unbuilt changes.

HTTP Status Code: 400

ConflictException

Two clients are using the same AWS account, Amazon Lex bot, and user ID.

HTTP Status Code: 409

InternalFailureException

Internal service error. Retry the call.

HTTP Status Code: 500

LimitExceededException

Exceeded a limit.

HTTP Status Code: 429

NotFoundException

The resource (such as the Amazon Lex bot or an alias) that is referred to is not found.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetSession

Service: Amazon Lex Runtime Service

Returns session information for a specified bot, alias, and user ID.

Request Syntax

```
GET /bot/botName/alias/botAlias/user/userId/session/?  
checkpointLabelFilter=checkpointLabelFilter HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

botAlias

The alias in use for the bot that contains the session data.

Required: Yes

botName

The name of the bot that contains the session data.

Required: Yes

checkpointLabelFilter

A string used to filter the intents returned in the `recentIntentSummaryView` structure.

When you specify a filter, only intents with their `checkpointLabel` field set to that string are returned.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[a-zA-Z0-9-]+`

userId

The ID of the client application user. Amazon Lex uses this to identify a user's conversation with your bot.

Length Constraints: Minimum length of 2. Maximum length of 100.

Pattern: `[0-9a-zA-Z._:-]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "activeContexts": [
    {
      "name": "string",
      "parameters": {
        "string" : "string"
      },
      "timeToLive": {
        "timeToLiveInSeconds": number,
        "turnsToLive": number
      }
    }
  ],
  "dialogAction": {
    "fulfillmentState": "string",
    "intentName": "string",
    "message": "string",
    "messageFormat": "string",
    "slots": {
      "string" : "string"
    },
    "slotToElicit": "string",
    "type": "string"
  },
  "recentIntentSummaryView": [
    {
      "checkpointLabel": "string",
      "confirmationStatus": "string",
      "dialogActionType": "string",
      "fulfillmentState": "string",
      "intentName": "string",
      "slots": {
```

```
        "string" : "string"
      },
      "slotToElicit": "string"
    }
  ],
  "sessionAttributes": {
    "string" : "string"
  },
  "sessionId": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

activeContexts

A list of active contexts for the session. A context can be set when an intent is fulfilled or by calling the `PostContent`, `PostText`, or `PutSession` operation.

You can use a context to control the intents that can follow up an intent, or to modify the operation of your application.

Type: Array of [ActiveContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 20 items.

dialogAction

Describes the current state of the bot.

Type: [DialogAction](#) object

recentIntentSummaryView

An array of information about the intents used in the session. The array can contain a maximum of three summaries. If more than three intents are used in the session, the `recentIntentSummaryView` operation contains information about the last three intents used.

If you set the `checkpointLabelFilter` parameter in the request, the array contains only the intents with the specified label.

Type: Array of [IntentSummary](#) objects

Array Members: Minimum number of 0 items. Maximum number of 3 items.

sessionAttributes

Map of key/value pairs representing the session-specific context information. It contains application information passed between Amazon Lex and a client application.

Type: String to string map

sessionId

A unique identifier for the session.

Type: String

Errors

BadRequestException

Request validation failed, there is no usable message in the context, or the bot build failed, is still in progress, or contains unbuilt changes.

HTTP Status Code: 400

InternalFailureException

Internal service error. Retry the call.

HTTP Status Code: 500

LimitExceededException

Exceeded a limit.

HTTP Status Code: 429

NotFoundException

The resource (such as the Amazon Lex bot or an alias) that is referred to is not found.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PostContent

Service: Amazon Lex Runtime Service

Sends user input (text or speech) to Amazon Lex. Clients use this API to send text and audio requests to Amazon Lex at runtime. Amazon Lex interprets the user input using the machine learning model that it built for the bot.

The PostContent operation supports audio input at 8kHz and 16kHz. You can use 8kHz audio to achieve higher speech recognition accuracy in telephone audio applications.

In response, Amazon Lex returns the next message to convey to the user. Consider the following example messages:

- For a user input "I would like a pizza," Amazon Lex might return a response with a message eliciting slot data (for example, PizzaSize): "What size pizza would you like?".
- After the user provides all of the pizza order information, Amazon Lex might return a response with a message to get user confirmation: "Order the pizza?".
- After the user replies "Yes" to the confirmation prompt, Amazon Lex might return a conclusion statement: "Thank you, your cheese pizza has been ordered.".

Not all Amazon Lex messages require a response from the user. For example, conclusion statements do not require a response. Some messages require only a yes or no response. In addition to the message, Amazon Lex provides additional context about the message in the response that you can use to enhance client behavior, such as displaying the appropriate client user interface. Consider the following examples:

- If the message is to elicit slot data, Amazon Lex returns the following context information:
 - x-amz-lex-dialog-state header set to ElicitSlot
 - x-amz-lex-intent-name header set to the intent name in the current context
 - x-amz-lex-slot-to-elicite header set to the slot name for which the message is eliciting information
 - x-amz-lex-slots header set to a map of slots configured for the intent with their current values
- If the message is a confirmation prompt, the x-amz-lex-dialog-state header is set to Confirmation and the x-amz-lex-slot-to-elicite header is omitted.

- If the message is a clarification prompt configured for the intent, indicating that the user intent is not understood, the `x-amz-dialog-state` header is set to `ElicitIntent` and the `x-amz-slot-to-elicit` header is omitted.

In addition, Amazon Lex also returns your application-specific `sessionAttributes`. For more information, see [Managing Conversation Context](#).

Request Syntax

```
POST /bot/botName/alias/botAlias/user/userId/content HTTP/1.1
x-amz-lex-session-attributes: sessionAttributes
x-amz-lex-request-attributes: requestAttributes
Content-Type: contentType
Accept: accept
x-amz-lex-active-contexts: activeContexts

inputStream
```

URI Request Parameters

The request uses the following URI parameters.

[accept](#)

You pass this value as the Accept HTTP header.

The message Amazon Lex returns in the response can be either text or speech based on the Accept HTTP header value in the request.

- If the value is `text/plain; charset=utf-8`, Amazon Lex returns text in the response.
- If the value begins with `audio/`, Amazon Lex returns speech in the response. Amazon Lex uses Amazon Polly to generate the speech (using the configuration you specified in the Accept header). For example, if you specify `audio/mpeg` as the value, Amazon Lex returns speech in the MPEG format.
- If the value is `audio/pcm`, the speech returned is `audio/pcm` in 16-bit, little endian format.
- The following are the accepted values:
 - `audio/mpeg`
 - `audio/ogg`
 - `audio/pcm`

- text/plain; charset=utf-8
- audio/* (defaults to mpeg)

activeContexts

A list of contexts active for the request. A context can be activated when a previous intent is fulfilled, or by including the context in the request,

If you don't specify a list of contexts, Amazon Lex will use the current list of contexts for the session. If you specify an empty list, all contexts for the session are cleared.

botAlias

Alias of the Amazon Lex bot.

Required: Yes

botName

Name of the Amazon Lex bot.

Required: Yes

contentType

You pass this value as the Content-Type HTTP header.

Indicates the audio format or text. The header value must start with one of the following prefixes:

- PCM format, audio data must be in little-endian byte order.
 - audio/l16; rate=16000; channels=1
 - audio/x-l16; sample-rate=16000; channel-count=1
 - audio/lpcm; sample-rate=8000; sample-size-bits=16; channel-count=1; is-big-endian=false
- Opus format
 - audio/x-cbr-opus-with-preamble; preamble-size=0; bit-rate=256000; frame-size-milliseconds=4
- Text format
 - text/plain; charset=utf-8

Required: Yes

requestAttributes

You pass this value as the `x-amz-lex-request-attributes` HTTP header.

Request-specific information passed between Amazon Lex and a client application. The value must be a JSON serialized and base64 encoded map with string keys and values. The total size of the `requestAttributes` and `sessionAttributes` headers is limited to 12 KB.

The namespace `x-amz-lex:` is reserved for special attributes. Don't create any request attributes with the prefix `x-amz-lex:`.

For more information, see [Setting Request Attributes](#).

sessionAttributes

You pass this value as the `x-amz-lex-session-attributes` HTTP header.

Application-specific information passed between Amazon Lex and a client application. The value must be a JSON serialized and base64 encoded map with string keys and values. The total size of the `sessionAttributes` and `requestAttributes` headers is limited to 12 KB.

For more information, see [Setting Session Attributes](#).

userId

The ID of the client application user. Amazon Lex uses this to identify a user's conversation with your bot. At runtime, each request must contain the `userId` field.

To decide the user ID to use for your application, consider the following factors.

- The `userId` field must not contain any personally identifiable information of the user, for example, name, personal identification numbers, or other end user personal information.
- If you want a user to start a conversation on one device and continue on another device, use a user-specific identifier.
- If you want the same user to be able to have two independent conversations on two different devices, choose a device-specific identifier.
- A user can't have two independent conversations with two different versions of the same bot. For example, a user can't have a conversation with the PROD and BETA versions of the same bot. If you anticipate that a user will need to have conversation with two different versions, for example, while testing, include the bot alias in the user ID to separate the two conversations.

Length Constraints: Minimum length of 2. Maximum length of 100.

Pattern: `[0-9a-zA-Z._:-]+`

Required: Yes

Request Body

The request accepts the following binary data.

inputStream

User input in PCM or Opus audio format or text format as described in the Content-Type HTTP header.

You can stream audio data to Amazon Lex or you can create a local buffer that captures all of the audio data before sending. In general, you get better performance if you stream audio data rather than buffering the data locally.

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-Type: contentType
x-amz-lex-intent-name: intentName
x-amz-lex-nlu-intent-confidence: nluIntentConfidence
x-amz-lex-alternative-intents: alternativeIntents
x-amz-lex-slots: slots
x-amz-lex-session-attributes: sessionAttributes
x-amz-lex-sentiment: sentimentResponse
x-amz-lex-message: message
x-amz-lex-encoded-message: encodedMessage
x-amz-lex-message-format: messageFormat
x-amz-lex-dialog-state: dialogState
x-amz-lex-slot-to-elicit: slotToElicit
x-amz-lex-input-transcript: inputTranscript
x-amz-lex-encoded-input-transcript: encodedInputTranscript
x-amz-lex-bot-version: botVersion
x-amz-lex-session-id: sessionId
x-amz-lex-active-contexts: activeContexts
```

audioStream

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

activeContexts

A list of active contexts for the session. A context can be set when an intent is fulfilled or by calling the `PostContent`, `PostText`, or `PutSession` operation.

You can use a context to control the intents that can follow up an intent, or to modify the operation of your application.

alternativeIntents

One to four alternative intents that may be applicable to the user's intent.

Each alternative includes a score that indicates how confident Amazon Lex is that the intent matches the user's intent. The intents are sorted by the confidence score.

botVersion

The version of the bot that responded to the conversation. You can use this information to help determine if one version of a bot is performing better than another version.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9]+|\$LATEST`

contentType

Content type as specified in the `Accept` HTTP header in the request.

dialogState

Identifies the current state of the user interaction. Amazon Lex returns one of the following values as `dialogState`. The client can optionally use this information to customize the user interface.

- `ElicitIntent` - Amazon Lex wants to elicit the user's intent. Consider the following examples:

For example, a user might utter an intent ("I want to order a pizza"). If Amazon Lex cannot infer the user intent from this utterance, it will return this dialog state.

- **ConfirmIntent** - Amazon Lex is expecting a "yes" or "no" response.

For example, Amazon Lex wants user confirmation before fulfilling an intent. Instead of a simple "yes" or "no" response, a user might respond with additional information. For example, "yes, but make it a thick crust pizza" or "no, I want to order a drink." Amazon Lex can process such additional information (in these examples, update the crust type slot or change the intent from OrderPizza to OrderDrink).

- **ElicitSlot** - Amazon Lex is expecting the value of a slot for the current intent.

For example, suppose that in the response Amazon Lex sends this message: "What size pizza would you like?". A user might reply with the slot value (e.g., "medium"). The user might also provide additional information in the response (e.g., "medium thick crust pizza"). Amazon Lex can process such additional information appropriately.

- **Fulfilled** - Conveys that the Lambda function has successfully fulfilled the intent.
- **ReadyForFulfillment** - Conveys that the client has to fulfill the request.
- **Failed** - Conveys that the conversation with the user failed.

This can happen for various reasons, including that the user does not provide an appropriate response to prompts from the service (you can configure how many times Amazon Lex can prompt a user for specific information), or if the Lambda function fails to fulfill the intent.

Valid Values: **ElicitIntent** | **ConfirmIntent** | **ElicitSlot** | **Fulfilled** | **ReadyForFulfillment** | **Failed**

encodedInputTranscript

The text used to process the request.

If the input was an audio stream, the `encodedInputTranscript` field contains the text extracted from the audio stream. This is the text that is actually processed to recognize intents and slot values. You can use this information to determine if Amazon Lex is correctly processing the audio that you send.

The `encodedInputTranscript` field is base-64 encoded. You must decode the field before you can use the value.

encodedMessage

The message to convey to the user. The message can come from the bot's configuration or from a Lambda function.

If the intent is not configured with a Lambda function, or if the Lambda function returned `Delegate` as the `dialogAction.type` in its response, Amazon Lex decides on the next course of action and selects an appropriate message from the bot's configuration based on the current interaction context. For example, if Amazon Lex isn't able to understand user input, it uses a clarification prompt message.

When you create an intent you can assign messages to groups. When messages are assigned to groups Amazon Lex returns one message from each group in the response. The message field is an escaped JSON string containing the messages. For more information about the structure of the JSON string returned, see [Supported Message Formats](#).

If the Lambda function returns a message, Amazon Lex passes it to the client in its response.

The `encodedMessage` field is base-64 encoded. You must decode the field before you can use the value.

Length Constraints: Minimum length of 1. Maximum length of 1366.

inputTranscript

This header has been deprecated.

The text used to process the request.

You can use this field only in the `de-DE`, `en-AU`, `en-GB`, `en-US`, `es-419`, `es-ES`, `es-US`, `fr-CA`, `fr-FR`, and `it-IT` locales. In all other locales, the `inputTranscript` field is null. You should use the `encodedInputTranscript` field instead.

If the input was an audio stream, the `inputTranscript` field contains the text extracted from the audio stream. This is the text that is actually processed to recognize intents and slot values. You can use this information to determine if Amazon Lex is correctly processing the audio that you send.

intentName

Current user intent that Amazon Lex is aware of.

message

This header has been deprecated.

You can only use this field in the de-DE, en-AU, en-GB, en-US, es-419, es-ES, es-US, fr-CA, fr-FR, and it-IT locales. In all other locales, the message field is null. You should use the `encodedMessage` field instead.

The message to convey to the user. The message can come from the bot's configuration or from a Lambda function.

If the intent is not configured with a Lambda function, or if the Lambda function returned `Delegate` as the `dialogAction.type` in its response, Amazon Lex decides on the next course of action and selects an appropriate message from the bot's configuration based on the current interaction context. For example, if Amazon Lex isn't able to understand user input, it uses a clarification prompt message.

When you create an intent you can assign messages to groups. When messages are assigned to groups Amazon Lex returns one message from each group in the response. The message field is an escaped JSON string containing the messages. For more information about the structure of the JSON string returned, see [Supported Message Formats](#).

If the Lambda function returns a message, Amazon Lex passes it to the client in its response.

Length Constraints: Minimum length of 1. Maximum length of 1024.

messageFormat

The format of the response message. One of the following values:

- `PlainText` - The message contains plain UTF-8 text.
- `CustomPayload` - The message is a custom format for the client.
- `SSML` - The message contains text formatted for voice output.
- `Composite` - The message contains an escaped JSON object containing one or more messages from the groups that messages were assigned to when the intent was created.

Valid Values: `PlainText` | `CustomPayload` | `SSML` | `Composite`

nlIntentConfidence

Provides a score that indicates how confident Amazon Lex is that the returned intent is the one that matches the user's intent. The score is between 0.0 and 1.0.

The score is a relative score, not an absolute score. The score may change based on improvements to Amazon Lex.

sentimentResponse

The sentiment expressed in an utterance.

When the bot is configured to send utterances to Amazon Comprehend for sentiment analysis, this field contains the result of the analysis.

sessionAttributes

Map of key/value pairs representing the session-specific context information.

sessionId

The unique identifier for the session.

slots

Map of zero or more intent slots (name/value pairs) Amazon Lex detected from the user input during the conversation. The field is base-64 encoded.

Amazon Lex creates a resolution list containing likely values for a slot. The value that it returns is determined by the `valueSelectionStrategy` selected when the slot type was created or updated. If `valueSelectionStrategy` is set to `ORIGINAL_VALUE`, the value provided by the user is returned, if the user value is similar to the slot values. If `valueSelectionStrategy` is set to `TOP_RESOLUTION` Amazon Lex returns the first value in the resolution list or, if there is no resolution list, null. If you don't specify a `valueSelectionStrategy`, the default is `ORIGINAL_VALUE`.

slotToElicit

If the `dialogState` value is `ElicitSlot`, returns the name of the slot for which Amazon Lex is eliciting a value.

The response returns the following as the HTTP body.

audioStream

The prompt (or statement) to convey to the user. This is based on the bot configuration and context. For example, if Amazon Lex did not understand the user intent, it sends the `clarificationPrompt` configured for the bot. If the intent requires confirmation before taking the fulfillment action, it sends the `confirmationPrompt`. Another example: Suppose

that the Lambda function successfully fulfilled the intent, and sent a message to convey to the user. Then Amazon Lex sends that message in the response.

Errors

BadGatewayException

Either the Amazon Lex bot is still building, or one of the dependent services (Amazon Polly, AWS Lambda) failed with an internal service error.

HTTP Status Code: 502

BadRequestException

Request validation failed, there is no usable message in the context, or the bot build failed, is still in progress, or contains unbuilt changes.

HTTP Status Code: 400

ConflictException

Two clients are using the same AWS account, Amazon Lex bot, and user ID.

HTTP Status Code: 409

DependencyFailedException

One of the dependencies, such as AWS Lambda or Amazon Polly, threw an exception. For example,

- If Amazon Lex does not have sufficient permissions to call a Lambda function.
- If a Lambda function takes longer than 30 seconds to execute.
- If a fulfillment Lambda function returns a `Delegat`e dialog action without removing any slot values.

HTTP Status Code: 424

InternalFailureException

Internal service error. Retry the call.

HTTP Status Code: 500

LimitExceededException

Exceeded a limit.

HTTP Status Code: 429

LoopDetectedException

This exception is not used.

HTTP Status Code: 508

NotAcceptableException

The accept header in the request does not have a valid value.

HTTP Status Code: 406

NotFoundException

The resource (such as the Amazon Lex bot or an alias) that is referred to is not found.

HTTP Status Code: 404

RequestTimeoutException

The input speech is too long.

HTTP Status Code: 408

UnsupportedMediaTypeException

The Content-Type header (PostContent API) has an invalid value.

HTTP Status Code: 415

Examples

Example 1

In this request, the URI identifies a bot (Traffic), bot version (\$LATEST), and end user name (someuser). The Content-Type header identifies the format of the audio in the body. Amazon Lex also supports other formats. To convert audio from one format to another, if necessary, you can use SoX open source software. You specify the format in which you want to get the response by adding the Accept HTTP header.

In the response, the x-amz-lex-message header shows the response that Amazon Lex returned. The client can then send this response to the user. The same message is sent in audio/MPEG format through chunked encoding (as requested).

Sample Request

```
"POST /bot/Traffic/alias/$LATEST/user/someuser/content HTTP/1.1[\r][\n]"
"x-amz-lex-session-attributes: eyJlc2VybmFtZSI6IkVvYiJ9[\r][\n]"
"Content-Type: audio/x-l16; channel-count=1; sample-rate=16000f[\r][\n]"
"Accept: audio/mpeg[\r][\n]"
"Host: runtime.lex.us-east-1.amazonaws.com[\r][\n]"
"Authorization: AWS4-HMAC-SHA256 Credential=BLANKED_OUT/20161230/us-east-1/lex/
aws4_request,
SignedHeaders=accept;content-type;host;x-amz-content-sha256;x-amz-date;x-amz-lex-
session-attributes,
Signature=78ca5b54ea3f64a17ff7522de02cd90a9acd2365b45a9ce9b96ea105bb1c7ec2[\r][\n]"
"X-Amz-Date: 20161230T181426Z[\r][\n]"
"X-Amz-Content-Sha256:
e3b0c44298fc1c149afbf4c8996fb92427ae41e4649b934ca495991b7852b855[\r][\n]"
"Transfer-Encoding: chunked[\r][\n]"
"Connection: Keep-Alive[\r][\n]"
"User-Agent: Apache-HttpClient/4.5.x (Java/1.8.0_112)[\r][\n]"
"Accept-Encoding: gzip,deflate[\r][\n]"
"[\r][\n]"
"1000[\r][\n]"
"[0x7][0x0][0x7][0x0][\n]"
"[0x0][0x7][0x0][0xfc][0xff][\n]"
"[0x0][\n]"
...
```

Sample Response

```
"HTTP/1.1 200 OK[\r][\n]"
"x-amzn-RequestId: cc8b34af-cebb-11e6-a35c-55f3a992f28d[\r][\n]"
"x-amz-lex-message: Sorry, can you repeat that?[\r][\n]"
"x-amz-lex-dialog-state: ElicitIntent[\r][\n]"
"x-amz-lex-session-attributes: eyJlc2VybmFtZSI6IkVvYiJ9[\r][\n]"
"Content-Type: audio/mpeg[\r][\n]"
"Transfer-Encoding: chunked[\r][\n]"
"Date: Fri, 30 Dec 2016 18:14:28 GMT[\r][\n]"
"[\r][\n]"
"2000[\r][\n]"
"ID3[0x4][0x0][0x0][0x0][0x0][0x0]#TSSE[0x0][0x0][0x0][0xf][0x0][0x0]
[0x3]Lavf57.41.100[0x0][0x0][0x0][0x0][0x0][0x0][0x0][0x0][0x0][0x0][0x0][0xff]
[0xf3]`[0xc4][0x0][0x1b]{[0x8d][0xe8][0x1]C[0x18][0x1][0x0]J[0xe0]`b[0xdd][0xd1]
[0xb][0xfd][0x11][0xdf][0xfe]";[0xbb][0xbb][0x9f][0xee][0xee][0xee][0xee]|DDD/[0xff]
[0xff][0xff][0xff]www?D[0xf7]w^?[0xff][0xfa]h[0x88][0x85][0xfe][0x88][0x88][0x88]"
```

```
[[0xa2]'[0xff][0xfa]"{[0x9f][0xe8][0x88]]D[0xeb][0xbb][0xbb][0xa2]!u[0xfd][0xdd][0xdf]
[0x88][0x94][0x0]F[0xef][0xa1]8[0x0][0x82]w[0x88]N[0x0][0x0][0x9b][0xbb][0xe8][0xe
...
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PostText

Service: Amazon Lex Runtime Service

Sends user input to Amazon Lex. Client applications can use this API to send requests to Amazon Lex at runtime. Amazon Lex then interprets the user input using the machine learning model it built for the bot.

In response, Amazon Lex returns the next message to convey to the user an optional `responseCard` to display. Consider the following example messages:

- For a user input "I would like a pizza", Amazon Lex might return a response with a message eliciting slot data (for example, `PizzaSize`): "What size pizza would you like?"
- After the user provides all of the pizza order information, Amazon Lex might return a response with a message to obtain user confirmation "Proceed with the pizza order?".
- After the user replies to a confirmation prompt with a "yes", Amazon Lex might return a conclusion statement: "Thank you, your cheese pizza has been ordered.".

Not all Amazon Lex messages require a user response. For example, a conclusion statement does not require a response. Some messages require only a "yes" or "no" user response. In addition to the message, Amazon Lex provides additional context about the message in the response that you might use to enhance client behavior, for example, to display the appropriate client user interface. These are the `slotToElicit`, `dialogState`, `intentName`, and `slots` fields in the response. Consider the following examples:

- If the message is to elicit slot data, Amazon Lex returns the following context information:
 - `dialogState` set to `ElicitSlot`
 - `intentName` set to the intent name in the current context
 - `slotToElicit` set to the slot name for which the message is eliciting information
 - `slots` set to a map of slots, configured for the intent, with currently known values
- If the message is a confirmation prompt, the `dialogState` is set to `ConfirmIntent` and `SlotToElicit` is set to null.
- If the message is a clarification prompt (configured for the intent) that indicates that user intent is not understood, the `dialogState` is set to `ElicitIntent` and `slotToElicit` is set to null.

In addition, Amazon Lex also returns your application-specific sessionAttributes. For more information, see [Managing Conversation Context](#).

Request Syntax

```
POST /bot/botName/alias/botAlias/user/userId/text HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "activeContexts": [
    {
      "name": "string",
      "parameters": {
        "string" : "string"
      },
      "timeToLive": {
        "timeToLiveInSeconds": number,
        "turnsToLive": number
      }
    }
  ],
  "inputText": "string",
  "requestAttributes": {
    "string" : "string"
  },
  "sessionAttributes": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

[botAlias](#)

The alias of the Amazon Lex bot.

Required: Yes

[botName](#)

The name of the Amazon Lex bot.

Required: Yes

userId

The ID of the client application user. Amazon Lex uses this to identify a user's conversation with your bot. At runtime, each request must contain the `userId` field.

To decide the user ID to use for your application, consider the following factors.

- The `userId` field must not contain any personally identifiable information of the user, for example, name, personal identification numbers, or other end user personal information.
- If you want a user to start a conversation on one device and continue on another device, use a user-specific identifier.
- If you want the same user to be able to have two independent conversations on two different devices, choose a device-specific identifier.
- A user can't have two independent conversations with two different versions of the same bot. For example, a user can't have a conversation with the PROD and BETA versions of the same bot. If you anticipate that a user will need to have conversation with two different versions, for example, while testing, include the bot alias in the user ID to separate the two conversations.

Length Constraints: Minimum length of 2. Maximum length of 100.

Pattern: `[0-9a-zA-Z._:-]+`

Required: Yes

Request Body

The request accepts the following data in JSON format.

activeContexts

A list of contexts active for the request. A context can be activated when a previous intent is fulfilled, or by including the context in the request,

If you don't specify a list of contexts, Amazon Lex will use the current list of contexts for the session. If you specify an empty list, all contexts for the session are cleared.

Type: Array of [ActiveContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 20 items.

Required: No

inputText

The text that the user entered (Amazon Lex interprets this text).

When you are using the AWS CLI, you can't pass a URL in the `--input-text` parameter. Pass the URL using the `--cli-input-json` parameter instead.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

requestAttributes

Request-specific information passed between Amazon Lex and a client application.

The namespace `x-amz-lex:` is reserved for special attributes. Don't create any request attributes with the prefix `x-amz-lex:`.

For more information, see [Setting Request Attributes](#).

Type: String to string map

Required: No

sessionAttributes

Application-specific information passed between Amazon Lex and a client application.

For more information, see [Setting Session Attributes](#).

Type: String to string map

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "activeContexts": [
    {
```

```

    "name": "string",
    "parameters": {
        "string" : "string"
    },
    "timeToLive": {
        "timeToLiveInSeconds": number,
        "turnsToLive": number
    }
},
"alternativeIntents": [
    {
        "intentName": "string",
        "nluIntentConfidence": {
            "score": number
        },
        "slots": {
            "string" : "string"
        }
    }
],
"botVersion": "string",
"dialogState": "string",
"intentName": "string",
"message": "string",
"messageFormat": "string",
"nluIntentConfidence": {
    "score": number
},
"responseCard": {
    "contentType": "string",
    "genericAttachments": [
        {
            "attachmentLinkUrl": "string",
            "buttons": [
                {
                    "text": "string",
                    "value": "string"
                }
            ],
            "imageUrl": "string",
            "subTitle": "string",
            "title": "string"
        }
    ]
}

```



```
    ],  
    "version": "string"  
  },  
  "sentimentResponse": {  
    "sentimentLabel": "string",  
    "sentimentScore": "string"  
  },  
  "sessionAttributes": {  
    "string" : "string"  
  },  
  "sessionId": "string",  
  "slots": {  
    "string" : "string"  
  },  
  "slotToElicit": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

activeContexts

A list of active contexts for the session. A context can be set when an intent is fulfilled or by calling the `PostContent`, `PostText`, or `PutSession` operation.

You can use a context to control the intents that can follow up an intent, or to modify the operation of your application.

Type: Array of [ActiveContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 20 items.

alternativeIntents

One to four alternative intents that may be applicable to the user's intent.

Each alternative includes a score that indicates how confident Amazon Lex is that the intent matches the user's intent. The intents are sorted by the confidence score.

Type: Array of [PredictedIntent](#) objects

Array Members: Maximum number of 4 items.

botVersion

The version of the bot that responded to the conversation. You can use this information to help determine if one version of a bot is performing better than another version.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9]+\|\\$LATEST`

dialogState

Identifies the current state of the user interaction. Amazon Lex returns one of the following values as `dialogState`. The client can optionally use this information to customize the user interface.

- `ElicitIntent` - Amazon Lex wants to elicit user intent.

For example, a user might utter an intent ("I want to order a pizza"). If Amazon Lex cannot infer the user intent from this utterance, it will return this `dialogState`.

- `ConfirmIntent` - Amazon Lex is expecting a "yes" or "no" response.

For example, Amazon Lex wants user confirmation before fulfilling an intent.

Instead of a simple "yes" or "no," a user might respond with additional information. For example, "yes, but make it thick crust pizza" or "no, I want to order a drink". Amazon Lex can process such additional information (in these examples, update the crust type slot value, or change intent from `OrderPizza` to `OrderDrink`).

- `ElicitSlot` - Amazon Lex is expecting a slot value for the current intent.

For example, suppose that in the response Amazon Lex sends this message: "What size pizza would you like?". A user might reply with the slot value (e.g., "medium"). The user might also provide additional information in the response (e.g., "medium thick crust pizza"). Amazon Lex can process such additional information appropriately.

- `Fulfilled` - Conveys that the Lambda function configured for the intent has successfully fulfilled the intent.
- `ReadyForFulfillment` - Conveys that the client has to fulfill the intent.
- `Failed` - Conveys that the conversation with the user failed.

This can happen for various reasons including that the user did not provide an appropriate response to prompts from the service (you can configure how many times Amazon Lex can prompt a user for specific information), or the Lambda function failed to fulfill the intent.

Type: String

Valid Values: `ElicitIntent` | `ConfirmIntent` | `ElicitSlot` | `Fulfilled` | `ReadyForFulfillment` | `Failed`

intentName

The current user intent that Amazon Lex is aware of.

Type: String

message

The message to convey to the user. The message can come from the bot's configuration or from a Lambda function.

If the intent is not configured with a Lambda function, or if the Lambda function returned `Delegate` as the `dialogAction.type` its response, Amazon Lex decides on the next course of action and selects an appropriate message from the bot's configuration based on the current interaction context. For example, if Amazon Lex isn't able to understand user input, it uses a clarification prompt message.

When you create an intent you can assign messages to groups. When messages are assigned to groups Amazon Lex returns one message from each group in the response. The message field is an escaped JSON string containing the messages. For more information about the structure of the JSON string returned, see [Supported Message Formats](#).

If the Lambda function returns a message, Amazon Lex passes it to the client in its response.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

messageFormat

The format of the response message. One of the following values:

- `PlainText` - The message contains plain UTF-8 text.
- `CustomPayload` - The message is a custom format defined by the Lambda function.

- SSML - The message contains text formatted for voice output.
- Composite - The message contains an escaped JSON object containing one or more messages from the groups that messages were assigned to when the intent was created.

Type: String

Valid Values: PlainText | CustomPayload | SSML | Composite

[nlIntentConfidence](#)

Provides a score that indicates how confident Amazon Lex is that the returned intent is the one that matches the user's intent. The score is between 0.0 and 1.0. For more information, see [Confidence Scores](#).

The score is a relative score, not an absolute score. The score may change based on improvements to Amazon Lex.

Type: [IntentConfidence](#) object

[responseCard](#)

Represents the options that the user has to respond to the current prompt. Response Card can come from the bot configuration (in the Amazon Lex console, choose the settings button next to a slot) or from a code hook (Lambda function).

Type: [ResponseCard](#) object

[sentimentResponse](#)

The sentiment expressed in an utterance.

When the bot is configured to send utterances to Amazon Comprehend for sentiment analysis, this field contains the result of the analysis.

Type: [SentimentResponse](#) object

[sessionAttributes](#)

A map of key-value pairs representing the session-specific context information.

Type: String to string map

[sessionId](#)

A unique identifier for the session.

Type: String

slots

The intent slots that Amazon Lex detected from the user input in the conversation.

Amazon Lex creates a resolution list containing likely values for a slot. The value that it returns is determined by the `valueSelectionStrategy` selected when the slot type was created or updated. If `valueSelectionStrategy` is set to `ORIGINAL_VALUE`, the value provided by the user is returned, if the user value is similar to the slot values. If `valueSelectionStrategy` is set to `TOP_RESOLUTION` Amazon Lex returns the first value in the resolution list or, if there is no resolution list, null. If you don't specify a `valueSelectionStrategy`, the default is `ORIGINAL_VALUE`.

Type: String to string map

slotToElicit

If the `dialogState` value is `ElicitSlot`, returns the name of the slot for which Amazon Lex is eliciting a value.

Type: String

Errors

BadGatewayException

Either the Amazon Lex bot is still building, or one of the dependent services (Amazon Polly, AWS Lambda) failed with an internal service error.

HTTP Status Code: 502

BadRequestException

Request validation failed, there is no usable message in the context, or the bot build failed, is still in progress, or contains unbuilt changes.

HTTP Status Code: 400

ConflictException

Two clients are using the same AWS account, Amazon Lex bot, and user ID.

HTTP Status Code: 409

DependencyFailedException

One of the dependencies, such as AWS Lambda or Amazon Polly, threw an exception. For example,

- If Amazon Lex does not have sufficient permissions to call a Lambda function.
- If a Lambda function takes longer than 30 seconds to execute.
- If a fulfillment Lambda function returns a De`l`egate dialog action without removing any slot values.

HTTP Status Code: 424

InternalFailureException

Internal service error. Retry the call.

HTTP Status Code: 500

LimitExceededException

Exceeded a limit.

HTTP Status Code: 429

LoopDetectedException

This exception is not used.

HTTP Status Code: 508

NotFoundException

The resource (such as the Amazon Lex bot or an alias) that is referred to is not found.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PutSession

Service: Amazon Lex Runtime Service

Creates a new session or modifies an existing session with an Amazon Lex bot. Use this operation to enable your application to set the state of the bot.

For more information, see [Managing Sessions](#).

Request Syntax

```
POST /bot/botName/alias/botAlias/user/userId/session HTTP/1.1
```

Accept: *accept*

Content-type: application/json

```
{
  "activeContexts": [
    {
      "name": "string",
      "parameters": {
        "string" : "string"
      },
      "timeToLive": {
        "timeToLiveInSeconds": number,
        "turnsToLive": number
      }
    }
  ],
  "dialogAction": {
    "fulfillmentState": "string",
    "intentName": "string",
    "message": "string",
    "messageFormat": "string",
    "slots": {
      "string" : "string"
    },
    "slotToElicit": "string",
    "type": "string"
  },
  "recentIntentSummaryView": [
    {
      "checkpointLabel": "string",
      "confirmationStatus": "string",
      "dialogActionType": "string",
```



```
    "fulfillmentState": "string",
    "intentName": "string",
    "slots": {
        "string" : "string"
    },
    "slotToElicit": "string"
}
],
"sessionAttributes": {
    "string" : "string"
}
}
```

URI Request Parameters

The request uses the following URI parameters.

accept

The message that Amazon Lex returns in the response can be either text or speech based depending on the value of this field.

- If the value is text/plain; charset=utf-8, Amazon Lex returns text in the response.
- If the value begins with audio/, Amazon Lex returns speech in the response. Amazon Lex uses Amazon Polly to generate the speech in the configuration that you specify. For example, if you specify audio/mpeg as the value, Amazon Lex returns speech in the MPEG format.
- If the value is audio/pcm, the speech is returned as audio/pcm in 16-bit, little endian format.
- The following are the accepted values:
 - audio/mpeg
 - audio/ogg
 - audio/pcm
 - audio/* (defaults to mpeg)
 - text/plain; charset=utf-8

botAlias

The alias in use for the bot that contains the session data.

Required: Yes

botName

The name of the bot that contains the session data.

Required: Yes

userId

The ID of the client application user. Amazon Lex uses this to identify a user's conversation with your bot.

Length Constraints: Minimum length of 2. Maximum length of 100.

Pattern: [0-9a-zA-Z._:-]+

Required: Yes

Request Body

The request accepts the following data in JSON format.

activeContexts

A list of contexts active for the request. A context can be activated when a previous intent is fulfilled, or by including the context in the request,

If you don't specify a list of contexts, Amazon Lex will use the current list of contexts for the session. If you specify an empty list, all contexts for the session are cleared.

Type: Array of [ActiveContext](#) objects

Array Members: Minimum number of 0 items. Maximum number of 20 items.

Required: No

dialogAction

Sets the next action that the bot should take to fulfill the conversation.

Type: [DialogAction](#) object

Required: No

[recentIntentSummaryView](#)

A summary of the recent intents for the bot. You can use the intent summary view to set a checkpoint label on an intent and modify attributes of intents. You can also use it to remove or add intent summary objects to the list.

An intent that you modify or add to the list must make sense for the bot. For example, the intent name must be valid for the bot. You must provide valid values for:

- `intentName`
- slot names
- `slotToElicit`

If you send the `recentIntentSummaryView` parameter in a `PutSession` request, the contents of the new summary view replaces the old summary view. For example, if a `GetSession` request returns three intents in the summary view and you call `PutSession` with one intent in the summary view, the next call to `GetSession` will only return one intent.

Type: Array of [IntentSummary](#) objects

Array Members: Minimum number of 0 items. Maximum number of 3 items.

Required: No

[sessionAttributes](#)

Map of key/value pairs representing the session-specific context information. It contains application information passed between Amazon Lex and a client application.

Type: String to string map

Required: No

Response Syntax

```
HTTP/1.1 200
Content-Type: contentType
x-amz-lex-intent-name: intentName
x-amz-lex-slots: slots
x-amz-lex-session-attributes: sessionAttributes
x-amz-lex-message: message
x-amz-lex-encoded-message: encodedMessage
x-amz-lex-message-format: messageFormat
```

```
x-amz-lex-dialog-state: dialogState  
x-amz-lex-slot-to-elicit: slotToElicit  
x-amz-lex-session-id: sessionId  
x-amz-lex-active-contexts: activeContexts
```

audioStream

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

activeContexts

A list of active contexts for the session.

contentType

Content type as specified in the Accept HTTP header in the request.

dialogState

- **ConfirmIntent** - Amazon Lex is expecting a "yes" or "no" response to confirm the intent before fulfilling an intent.
- **ElicitIntent** - Amazon Lex wants to elicit the user's intent.
- **ElicitSlot** - Amazon Lex is expecting the value of a slot for the current intent.
- **Failed** - Conveys that the conversation with the user has failed. This can happen for various reasons, including the user does not provide an appropriate response to prompts from the service, or if the Lambda function fails to fulfill the intent.
- **Fulfilled** - Conveys that the Lambda function has successfully fulfilled the intent.
- **ReadyForFulfillment** - Conveys that the client has to fulfill the intent.

Valid Values: ElicitIntent | ConfirmIntent | ElicitSlot | Fulfilled | ReadyForFulfillment | Failed

encodedMessage

The next message that should be presented to the user.

The encodedMessage field is base-64 encoded. You must decode the field before you can use the value.

Length Constraints: Minimum length of 1. Maximum length of 1366.

intentName

The name of the current intent.

message

This header has been deprecated.

The next message that should be presented to the user.

You can only use this field in the de-DE, en-AU, en-GB, en-US, es-419, es-ES, es-US, fr-CA, fr-FR, and it-IT locales. In all other locales, the message field is null. You should use the encodedMessage field instead.

Length Constraints: Minimum length of 1. Maximum length of 1024.

messageFormat

The format of the response message. One of the following values:

- PlainText - The message contains plain UTF-8 text.
- CustomPayload - The message is a custom format for the client.
- SSML - The message contains text formatted for voice output.
- Composite - The message contains an escaped JSON object containing one or more messages from the groups that messages were assigned to when the intent was created.

Valid Values: PlainText | CustomPayload | SSML | Composite

sessionAttributes

Map of key/value pairs representing session-specific context information.

sessionId

A unique identifier for the session.

slots

Map of zero or more intent slots Amazon Lex detected from the user input during the conversation.

Amazon Lex creates a resolution list containing likely values for a slot. The value that it returns is determined by the valueSelectionStrategy selected when the slot type was created or

updated. If `valueSelectionStrategy` is set to `ORIGINAL_VALUE`, the value provided by the user is returned, if the user value is similar to the slot values. If `valueSelectionStrategy` is set to `TOP_RESOLUTION` Amazon Lex returns the first value in the resolution list or, if there is no resolution list, null. If you don't specify a `valueSelectionStrategy` the default is `ORIGINAL_VALUE`.

slotToElicit

If the `dialogState` is `ElicitSlot`, returns the name of the slot for which Amazon Lex is eliciting a value.

The response returns the following as the HTTP body.

audioStream

The audio version of the message to convey to the user.

Errors

BadGatewayException

Either the Amazon Lex bot is still building, or one of the dependent services (Amazon Polly, AWS Lambda) failed with an internal service error.

HTTP Status Code: 502

BadRequestException

Request validation failed, there is no usable message in the context, or the bot build failed, is still in progress, or contains unbuilt changes.

HTTP Status Code: 400

ConflictException

Two clients are using the same AWS account, Amazon Lex bot, and user ID.

HTTP Status Code: 409

DependencyFailedException

One of the dependencies, such as AWS Lambda or Amazon Polly, threw an exception. For example,

- If Amazon Lex does not have sufficient permissions to call a Lambda function.
- If a Lambda function takes longer than 30 seconds to execute.
- If a fulfillment Lambda function returns a De`l`egate dialog action without removing any slot values.

HTTP Status Code: 424

InternalFailureException

Internal service error. Retry the call.

HTTP Status Code: 500

LimitExceededException

Exceeded a limit.

HTTP Status Code: 429

NotAcceptableException

The accept header in the request does not have a valid value.

HTTP Status Code: 406

NotFoundException

The resource (such as the Amazon Lex bot or an alias) that is referred to is not found.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

Data Types

The following data types are supported by Amazon Lex Model Building Service:

- [BotAliasMetadata](#)
- [BotChannelAssociation](#)
- [BotMetadata](#)
- [BuiltinIntentMetadata](#)
- [BuiltinIntentSlot](#)
- [BuiltinSlotTypeMetadata](#)
- [CodeHook](#)
- [ConversationLogsRequest](#)
- [ConversationLogsResponse](#)
- [EnumerationValue](#)
- [FollowUpPrompt](#)
- [FulfillmentActivity](#)
- [InputContext](#)
- [Intent](#)
- [IntentMetadata](#)
- [KendraConfiguration](#)
- [LogSettingsRequest](#)
- [LogSettingsResponse](#)
- [Message](#)
- [MigrationAlert](#)
- [MigrationSummary](#)
- [OutputContext](#)

- [Prompt](#)
- [ResourceReference](#)
- [Slot](#)
- [SlotDefaultValue](#)
- [SlotDefaultValueSpec](#)
- [SlotTypeConfiguration](#)
- [SlotTypeMetadata](#)
- [SlotTypeRegexConfiguration](#)
- [Statement](#)
- [Tag](#)
- [UtteranceData](#)
- [UtteranceList](#)

The following data types are supported by Amazon Lex Runtime Service:

- [ActiveContext](#)
- [ActiveContextTimeToLive](#)
- [Button](#)
- [DialogAction](#)
- [GenericAttachment](#)
- [IntentConfidence](#)
- [IntentSummary](#)
- [PredictedIntent](#)
- [ResponseCard](#)
- [SentimentResponse](#)

Amazon Lex Model Building Service

The following data types are supported by Amazon Lex Model Building Service:

- [BotAliasMetadata](#)
- [BotChannelAssociation](#)

- [BotMetadata](#)
- [BuiltinIntentMetadata](#)
- [BuiltinIntentSlot](#)
- [BuiltinSlotTypeMetadata](#)
- [CodeHook](#)
- [ConversationLogsRequest](#)
- [ConversationLogsResponse](#)
- [EnumerationValue](#)
- [FollowUpPrompt](#)
- [FulfillmentActivity](#)
- [InputContext](#)
- [Intent](#)
- [IntentMetadata](#)
- [KendraConfiguration](#)
- [LogSettingsRequest](#)
- [LogSettingsResponse](#)
- [Message](#)
- [MigrationAlert](#)
- [MigrationSummary](#)
- [OutputContext](#)
- [Prompt](#)
- [ResourceReference](#)
- [Slot](#)
- [SlotDefaultValue](#)
- [SlotDefaultValueSpec](#)
- [SlotTypeConfiguration](#)
- [SlotTypeMetadata](#)
- [SlotTypeRegexConfiguration](#)
- [Statement](#)
- [Tag](#)

- [UtteranceData](#)
- [UtteranceList](#)

BotAliasMetadata

Service: Amazon Lex Model Building Service

Provides information about a bot alias.

Contents

botName

The name of the bot to which the alias points.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: No

botVersion

The version of the Amazon Lex bot to which the alias points.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: No

checksum

Checksum of the bot alias.

Type: String

Required: No

conversationLogs

Settings that determine how Amazon Lex uses conversation logs for the alias.

Type: [ConversationLogsResponse](#) object

Required: No

createdDate

The date that the bot alias was created.

Type: Timestamp

Required: No

description

A description of the bot alias.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

lastUpdatedDate

The date that the bot alias was updated. When you create a resource, the creation date and last updated date are the same.

Type: Timestamp

Required: No

name

The name of the bot alias.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BotChannelAssociation

Service: Amazon Lex Model Building Service

Represents an association between an Amazon Lex bot and an external messaging platform.

Contents

botAlias

An alias pointing to the specific version of the Amazon Lex bot to which this association is being made.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: No

botConfiguration

Provides information necessary to communicate with the messaging platform.

Type: String to string map

Map Entries: Maximum number of 10 items.

Required: No

botName

The name of the Amazon Lex bot to which this association is being made.

Note

Currently, Amazon Lex supports associations with Facebook and Slack, and Twilio.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: `^[A-Za-z_?]+$`

Required: No

createdDate

The date that the association between the Amazon Lex bot and the channel was created.

Type: Timestamp

Required: No

description

A text description of the association you are creating.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

failureReason

If status is FAILED, Amazon Lex provides the reason that it failed to create the association.

Type: String

Required: No

name

The name of the association between the bot and the channel.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: No

status

The status of the bot channel.

- **CREATED** - The channel has been created and is ready for use.
- **IN_PROGRESS** - Channel creation is in progress.

- FAILED - There was an error creating the channel. For information about the reason for the failure, see the `failureReason` field.

Type: String

Valid Values: IN_PROGRESS | CREATED | FAILED

Required: No

type

Specifies the type of association by indicating the type of channel being established between the Amazon Lex bot and the external messaging platform.

Type: String

Valid Values: Facebook | Slack | Twilio-Sms | Kik

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BotMetadata

Service: Amazon Lex Model Building Service

Provides information about a bot. .

Contents

createdDate

The date that the bot was created.

Type: Timestamp

Required: No

description

A description of the bot.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

lastUpdatedDate

The date that the bot was updated. When you create a bot, the creation date and last updated date are the same.

Type: Timestamp

Required: No

name

The name of the bot.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z_?)+\$

Required: No

status

The status of the bot.

Type: String

Valid Values: BUILDING | READY | READY_BASIC_TESTING | FAILED | NOT_BUILT

Required: No

version

The version of the bot. For a new bot, the version is always \$LATEST.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: \ \$LATEST | [0-9] +

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BuiltinIntentMetadata

Service: Amazon Lex Model Building Service

Provides metadata for a built-in intent.

Contents

signature

A unique identifier for the built-in intent. To find the signature for an intent, see [Standard Built-in Intents](#) in the *Alexa Skills Kit*.

Type: String

Required: No

supportedLocales

A list of identifiers for the locales that the intent supports.

Type: Array of strings

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BuiltinIntentSlot

Service: Amazon Lex Model Building Service

Provides information about a slot used in a built-in intent.

Contents

name

A list of the slots defined for the intent.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BuiltinSlotTypeMetadata

Service: Amazon Lex Model Building Service

Provides information about a built in slot type.

Contents

signature

A unique identifier for the built-in slot type. To find the signature for a slot type, see [Slot Type Reference](#) in the *Alexa Skills Kit*.

Type: String

Required: No

supportedLocales

A list of target locales for the slot.

Type: Array of strings

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CodeHook

Service: Amazon Lex Model Building Service

Specifies a Lambda function that verifies requests to a bot or fulfills the user's request to a bot..

Contents

messageVersion

The version of the request-response that you want Amazon Lex to use to invoke your Lambda function. For more information, see [Using Lambda Functions](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 5.

Required: Yes

uri

The Amazon Resource Name (ARN) of the Lambda function.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:aws[a-zA-Z-]*:lambda:[a-z]+-[a-z]+(-[a-z]+)*-[0-9]:[0-9]{12}:function:[a-zA-Z0-9-_\u002F][0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})?(:[a-zA-Z0-9-_\u002F]+)?`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ConversationLogsRequest

Service: Amazon Lex Model Building Service

Provides the settings needed for conversation logs.

Contents

iamRoleArn

The Amazon Resource Name (ARN) of an IAM role with permission to write to your CloudWatch Logs for text logs and your S3 bucket for audio logs. If audio encryption is enabled, this role also provides access permission for the AWS KMS key used for encrypting audio logs. For more information, see [Creating an IAM Role and Policy for Conversation Logs](#).

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `^arn:[\w\ -]+:iam:[\d]{12}:role/.+$`

Required: Yes

logSettings

The settings for your conversation logs. You can log the conversation text, conversation audio, or both.

Type: Array of [LogSettingsRequest](#) objects

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ConversationLogsResponse

Service: Amazon Lex Model Building Service

Contains information about conversation log settings.

Contents

iamRoleArn

The Amazon Resource Name (ARN) of the IAM role used to write your logs to CloudWatch Logs or an S3 bucket.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `^arn:[\w\ -]+:iam:[\d]{12}:role/.+$`

Required: No

logSettings

The settings for your conversation logs. You can log text, audio, or both.

Type: Array of [LogSettingsResponse](#) objects

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

EnumerationValue

Service: Amazon Lex Model Building Service

Each slot type can have a set of values. Each enumeration value represents a value the slot type can take.

For example, a pizza ordering bot could have a slot type that specifies the type of crust that the pizza should have. The slot type could include the values

- thick
- thin
- stuffed

Contents

value

The value of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 140.

Required: Yes

synonyms

Additional values related to the slot type value.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 140.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

FollowUpPrompt

Service: Amazon Lex Model Building Service

A prompt for additional activity after an intent is fulfilled. For example, after the `OrderPizza` intent is fulfilled, you might prompt the user to find out whether the user wants to order drinks.

Contents

prompt

Prompts for information from the user.

Type: [Prompt](#) object

Required: Yes

rejectionStatement

If the user answers "no" to the question defined in the prompt field, Amazon Lex responds with this statement to acknowledge that the intent was canceled.

Type: [Statement](#) object

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

FulfillmentActivity

Service: Amazon Lex Model Building Service

Describes how the intent is fulfilled after the user provides all of the information required for the intent. You can provide a Lambda function to process the intent, or you can return the intent information to the client application. We recommend that you use a Lambda function so that the relevant logic lives in the Cloud and limit the client-side code primarily to presentation. If you need to update the logic, you only update the Lambda function; you don't need to upgrade your client application.

Consider the following examples:

- In a pizza ordering application, after the user provides all of the information for placing an order, you use a Lambda function to place an order with a pizzeria.
- In a gaming application, when a user says "pick up a rock," this information must go back to the client application so that it can perform the operation and update the graphics. In this case, you want Amazon Lex to return the intent data to the client.

Contents

type

How the intent should be fulfilled, either by running a Lambda function or by returning the slot data to the client application.

Type: String

Valid Values: ReturnIntent | CodeHook

Required: Yes

codeHook

A description of the Lambda function that is run to fulfill the intent.

Type: [CodeHook](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InputContext

Service: Amazon Lex Model Building Service

The name of a context that must be active for an intent to be selected by Amazon Lex.

Contents

name

The name of the context.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Intent

Service: Amazon Lex Model Building Service

Identifies the specific version of an intent.

Contents

intentName

The name of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

intentVersion

The version of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IntentMetadata

Service: Amazon Lex Model Building Service

Provides information about an intent.

Contents

createdDate

The date that the intent was created.

Type: Timestamp

Required: No

description

A description of the intent.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

lastUpdatedDate

The date that the intent was updated. When you create an intent, the creation date and last updated date are the same.

Type: Timestamp

Required: No

name

The name of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z_?])+`

Required: No

version

The version of the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

KendraConfiguration

Service: Amazon Lex Model Building Service

Provides configuration information for the AMAZON.KendraSearchIntent intent. When you use this intent, Amazon Lex searches the specified Amazon Kendra index and returns documents from the index that match the user's utterance. For more information, see [AMAZON.KendraSearchIntent](#).

Contents

kendraIndex

The Amazon Resource Name (ARN) of the Amazon Kendra index that you want the AMAZON.KendraSearchIntent intent to search. The index must be in the same account and Region as the Amazon Lex bot. If the Amazon Kendra index does not exist, you get an exception when you call the PutIntent operation.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:aws:kendra:[a-z]+-[a-z]+-[0-9]:[0-9]{12}:index\[a-zA-Z0-9\][a-zA-Z0-9_-]*`

Required: Yes

role

The Amazon Resource Name (ARN) of an IAM role that has permission to search the Amazon Kendra index. The role must be in the same account and Region as the Amazon Lex bot. If the role does not exist, you get an exception when you call the PutIntent operation.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:aws:iam::[0-9]{12}:role/.*`

Required: Yes

queryFilterString

A query filter that Amazon Lex sends to Amazon Kendra to filter the response from the query. The filter is in the format defined by Amazon Kendra. For more information, see [Filtering queries](#).

You can override this filter string with a new filter string at runtime.

Type: String

Length Constraints: Minimum length of 0.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LogSettingsRequest

Service: Amazon Lex Model Building Service

Settings used to configure delivery mode and destination for conversation logs.

Contents

destination

Where the logs will be delivered. Text logs are delivered to a CloudWatch Logs log group. Audio logs are delivered to an S3 bucket.

Type: String

Valid Values: CLOUDWATCH_LOGS | S3

Required: Yes

logType

The type of logging to enable. Text logs are delivered to a CloudWatch Logs log group. Audio logs are delivered to an S3 bucket.

Type: String

Valid Values: AUDIO | TEXT

Required: Yes

resourceArn

The Amazon Resource Name (ARN) of the CloudWatch Logs log group or S3 bucket where the logs should be delivered.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `^arn:[\w\-\+](?:logs:[\w\-\+]:[\d]{12}:log-group:[\.\-_\/#A-Za-z0-9]{1,512}(?::\*?)?|s3:::[a-z0-9][\.\-a-z0-9]{1,61}[a-z0-9])$`

Required: Yes

kmsKeyArn

The Amazon Resource Name (ARN) of the AWS KMS customer managed key for encrypting audio logs delivered to an S3 bucket. The key does not apply to CloudWatch Logs and is optional for S3 buckets.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `^arn:[\w\-\-]+:kms:[\w\-\-]+:[\d]{12}:(?:key\/[\w\-\-]+|alias\/[a-zA-Z0-9:\_\-\-]{1,256})$`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LogSettingsResponse

Service: Amazon Lex Model Building Service

The settings for conversation logs.

Contents

destination

The destination where logs are delivered.

Type: String

Valid Values: CLOUDWATCH_LOGS | S3

Required: No

kmsKeyArn

The Amazon Resource Name (ARN) of the key used to encrypt audio logs in an S3 bucket.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `^arn:[\w\-\+]:kms:[\w\-\+]:[\d]{12}:(?:key\/[\w\-\+]|alias\/[a-zA-Z0-9:\_\-\-]{1,256})$`

Required: No

logType

The type of logging that is enabled.

Type: String

Valid Values: AUDIO | TEXT

Required: No

resourceArn

The Amazon Resource Name (ARN) of the CloudWatch Logs log group or S3 bucket where the logs are delivered.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `^arn:[\w\-\+](?:logs:[\w\-\+]:[\d]{12}:log-group:[\.\-\_/#A-Za-z0-9]{1,512}(?::\*)?|s3:::[a-z0-9][\.\-a-z0-9]{1,61}[a-z0-9])$`

Required: No

resourcePrefix

The resource prefix is the first part of the S3 object key within the S3 bucket that you specified to contain audio logs. For CloudWatch Logs it is the prefix of the log stream name within the log group that you specified.

Type: String

Length Constraints: Maximum length of 1024.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Message

Service: Amazon Lex Model Building Service

The message object that provides the message text and its type.

Contents

content

The text of the message.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1000.

Required: Yes

contentType

The content type of the message string.

Type: String

Valid Values: PlainText | SSML | CustomPayload

Required: Yes

groupNumber

Identifies the message group that the message belongs to. When a group is assigned to a message, Amazon Lex returns one message from each group in the response.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 5.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MigrationAlert

Service: Amazon Lex Model Building Service

Provides information about alerts and warnings that Amazon Lex sends during a migration. The alerts include information about how to resolve the issue.

Contents

details

Additional details about the alert.

Type: Array of strings

Required: No

message

A message that describes why the alert was issued.

Type: String

Required: No

referenceURLs

A link to the Amazon Lex documentation that describes how to resolve the alert.

Type: Array of strings

Required: No

type

The type of alert. There are two kinds of alerts:

- **ERROR** - There was an issue with the migration that can't be resolved. The migration stops.
- **WARN** - There was an issue with the migration that requires manual changes to the new Amazon Lex V2 bot. The migration continues.

Type: String

Valid Values: **ERROR** | **WARN**

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MigrationSummary

Service: Amazon Lex Model Building Service

Provides information about migrating a bot from Amazon Lex V1 to Amazon Lex V2.

Contents

migrationId

The unique identifier that Amazon Lex assigned to the migration.

Type: String

Length Constraints: Fixed length of 10.

Pattern: `^[0-9a-zA-Z]+$`

Required: No

migrationStatus

The status of the operation. When the status is COMPLETE the bot is available in Amazon Lex V2. There may be alerts and warnings that need to be resolved to complete the migration.

Type: String

Valid Values: IN_PROGRESS | COMPLETED | FAILED

Required: No

migrationStrategy

The strategy used to conduct the migration.

Type: String

Valid Values: CREATE_NEW | UPDATE_EXISTING

Required: No

migrationTimestamp

The date and time that the migration started.

Type: Timestamp

Required: No

v1BotLocale

The locale of the Amazon Lex V1 bot that is the source of the migration.

Type: String

Valid Values: de-DE | en-AU | en-GB | en-IN | en-US | es-419 | es-ES | es-US
| fr-FR | fr-CA | it-IT | ja-JP | ko-KR

Required: No

v1BotName

The name of the Amazon Lex V1 bot that is the source of the migration.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 50.

Pattern: ^([A-Za-z]_?)+\$

Required: No

v1BotVersion

The version of the Amazon Lex V1 bot that is the source of the migration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: \\$(LATEST|[0-9])+

Required: No

v2BotId

The unique identifier of the Amazon Lex V2 that is the destination of the migration.

Type: String

Length Constraints: Fixed length of 10.

Pattern: ^[0-9a-zA-Z]+\$

Required: No

v2BotRole

The IAM role that Amazon Lex uses to run the Amazon Lex V2 bot.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `^arn:[\w\ -]+:iam::[\d]{12}:role/.[+ ]$`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

OutputContext

Service: Amazon Lex Model Building Service

The specification of an output context that is set when an intent is fulfilled.

Contents

name

The name of the context.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

timeToLiveInSeconds

The number of seconds that the context should be active after it is first sent in a `PostContent` or `PostText` response. You can set the value between 5 and 86,400 seconds (24 hours).

Type: Integer

Valid Range: Minimum value of 5. Maximum value of 86400.

Required: Yes

turnsToLive

The number of conversation turns that the context should be active. A conversation turn is one `PostContent` or `PostText` request and the corresponding response from Amazon Lex.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 20.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Prompt

Service: Amazon Lex Model Building Service

Obtains information from the user. To define a prompt, provide one or more messages and specify the number of attempts to get information from the user. If you provide more than one message, Amazon Lex chooses one of the messages to use to prompt the user. For more information, see [Amazon Lex: How It Works](#).

Contents

maxAttempts

The number of times to prompt the user for information.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 5.

Required: Yes

messages

An array of objects, each of which provides a message string and its type. You can specify the message string in plain text or in Speech Synthesis Markup Language (SSML).

Type: Array of [Message](#) objects

Array Members: Minimum number of 1 item. Maximum number of 15 items.

Required: Yes

responseCard

A response card. Amazon Lex uses this prompt at runtime, in the PostText API response. It substitutes session attributes and slot values for placeholders in the response card. For more information, see [Using a Response Card](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 50000.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ResourceReference

Service: Amazon Lex Model Building Service

Describes the resource that refers to the resource that you are attempting to delete. This object is returned as part of the `ResourceInUseException` exception.

Contents

name

The name of the resource that is using the resource that you are trying to delete.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `[a-zA-Z_]+`

Required: No

version

The version of the resource that is using the resource that you are trying to delete.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Slot

Service: Amazon Lex Model Building Service

Identifies the version of a specific slot.

Contents

name

The name of the slot.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z](-|_|. )?)+$`

Required: Yes

slotConstraint

Specifies whether the slot is required or optional.

Type: String

Valid Values: Required | Optional

Required: Yes

defaultValueSpec

A list of default values for the slot. Default values are used when Amazon Lex hasn't determined a value for a slot. You can specify default values from context variables, session attributes, and defined values.

Type: [SlotDefaultValueSpec](#) object

Required: No

description

A description of the slot.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

obfuscationSetting

Determines whether a slot is obfuscated in conversation logs and stored utterances. When you obfuscate a slot, the value is replaced by the slot name in curly braces ({}). For example, if the slot name is "full_name", obfuscated values are replaced with "{full_name}". For more information, see [Slot Obfuscation](#).

Type: String

Valid Values: NONE | DEFAULT_OBFUSCATION

Required: No

priority

Directs Amazon Lex the order in which to elicit this slot value from the user. For example, if the intent has two slots with priorities 1 and 2, AWS Amazon Lex first elicits a value for the slot with priority 1.

If multiple slots share the same priority, the order in which Amazon Lex elicits values is arbitrary.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

responseCard

A set of possible responses for the slot type used by text-based clients. A user chooses an option from the response card, instead of using text to reply.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 50000.

Required: No

sampleUtterances

If you know a specific pattern with which users might respond to an Amazon Lex request for a slot value, you can provide those utterances to improve accuracy. This is optional. In most cases, Amazon Lex is capable of understanding user utterances.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Length Constraints: Minimum length of 1. Maximum length of 200.

Required: No

slotType

The type of the slot, either a custom slot type that you defined or one of the built-in slot types.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^((AMAZON\.)_?|[A-Za-z]_?)+`

Required: No

slotTypeVersion

The version of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: No

valueElicitationPrompt

The prompt that Amazon Lex uses to elicit the slot value from the user.

Type: [Prompt](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SlotDefaultValue

Service: Amazon Lex Model Building Service

A default value for a slot.

Contents

defaultValue

The default value for the slot. You can specify one of the following:

- `#context-name.slot-name` - The slot value "slot-name" in the context "context-name."
- `{attribute}` - The slot value of the session attribute "attribute."
- `'value'` - The discrete value "value."

Type: String

Length Constraints: Minimum length of 1. Maximum length of 202.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SlotDefaultValueSpec

Service: Amazon Lex Model Building Service

Contains the default values for a slot. Default values are used when Amazon Lex hasn't determined a value for a slot.

Contents

defaultValueList

The default values for a slot. You can specify more than one default. For example, you can specify a default value to use from a matching context variable, a session attribute, or a fixed value.

The default value chosen is selected based on the order that you specify them in the list. For example, if you specify a context variable and a fixed value in that order, Amazon Lex uses the context variable if it is available, else it uses the fixed value.

Type: Array of [SlotDefaultValue](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SlotTypeConfiguration

Service: Amazon Lex Model Building Service

Provides configuration information for a slot type.

Contents

regexConfiguration

A regular expression used to validate the value of a slot.

Type: [SlotTypeRegexConfiguration](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SlotTypeMetadata

Service: Amazon Lex Model Building Service

Provides information about a slot type..

Contents

createdDate

The date that the slot type was created.

Type: Timestamp

Required: No

description

A description of the slot type.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 200.

Required: No

lastUpdatedDate

The date that the slot type was updated. When you create a resource, the creation date and last updated date are the same.

Type: Timestamp

Required: No

name

The name of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^([A-Za-z]_?)+$`

Required: No

version

The version of the slot type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SlotTypeRegexConfiguration

Service: Amazon Lex Model Building Service

Provides a regular expression used to validate the value of a slot.

Contents

pattern

A regular expression used to validate the value of a slot.

Use a standard regular expression. Amazon Lex supports the following characters in the regular expression:

- A-Z, a-z
- 0-9
- Unicode characters ("`\ u<Unicode>`")

Represent Unicode characters with four digits, for example "`\u0041`" or "`\u005A`".

The following regular expression operators are not supported:

- Infinite repeaters: `*`, `+`, or `{x,}` with no upper bound.
- Wild card (`.`)

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Statement

Service: Amazon Lex Model Building Service

A collection of messages that convey information to the user. At runtime, Amazon Lex selects the message to convey.

Contents

messages

A collection of message objects.

Type: Array of [Message](#) objects

Array Members: Minimum number of 1 item. Maximum number of 15 items.

Required: Yes

responseCard

At runtime, if the client is using the [PostText](#) API, Amazon Lex includes the response card in the response. It substitutes all of the session attributes and slot values for placeholders in the response card.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 50000.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Tag

Service: Amazon Lex Model Building Service

A list of key/value pairs that identify a bot, bot alias, or bot channel. Tag keys and values can consist of Unicode letters, digits, white space, and any of the following symbols: _ . : / = + - @.

Contents

key

The key for the tag. Keys are not case-sensitive and must be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

value

The value associated with a key. The value may be an empty string but it can't be null.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

UtteranceData

Service: Amazon Lex Model Building Service

Provides information about a single utterance that was made to your bot.

Contents

count

The number of times that the utterance was processed.

Type: Integer

Required: No

distinctUsers

The total number of individuals that used the utterance.

Type: Integer

Required: No

firstUtteredDate

The date that the utterance was first recorded.

Type: Timestamp

Required: No

lastUtteredDate

The date that the utterance was last recorded.

Type: Timestamp

Required: No

utteranceString

The text that was entered by the user or the text representation of an audio clip.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

UtteranceList

Service: Amazon Lex Model Building Service

Provides a list of utterances that have been made to a specific version of your bot. The list contains a maximum of 100 utterances.

Contents

botVersion

The version of the bot that processed the list.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `\$LATEST|[0-9]+`

Required: No

utterances

One or more [UtteranceData](#) objects that contain information about the utterances that have been made to a bot. The maximum number of object is 100.

Type: Array of [UtteranceData](#) objects

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Amazon Lex Runtime Service

The following data types are supported by Amazon Lex Runtime Service:

- [ActiveContext](#)
- [ActiveContextTimeToLive](#)
- [Button](#)
- [DialogAction](#)
- [GenericAttachment](#)
- [IntentConfidence](#)
- [IntentSummary](#)
- [PredictedIntent](#)
- [ResponseCard](#)
- [SentimentResponse](#)

ActiveContext

Service: Amazon Lex Runtime Service

A context is a variable that contains information about the current state of the conversation between a user and Amazon Lex. Context can be set automatically by Amazon Lex when an intent is fulfilled, or it can be set at runtime using the `PutContent`, `PutText`, or `PutSession` operation.

Contents

name

The name of the context.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Pattern: `^[A-Za-z_?]+$`

Required: Yes

parameters

State variables for the current context. You can use these values as default values for slots in subsequent events.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 10 items.

Key Length Constraints: Minimum length of 1. Maximum length of 100.

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

timeToLive

The length of time or number of turns that a context remains active.

Type: [ActiveContextTimeToLive](#) object

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ActiveContextTimeToLive

Service: Amazon Lex Runtime Service

The length of time or number of turns that a context remains active.

Contents

timeToLiveInSeconds

The number of seconds that the context should be active after it is first sent in a `PostContent` or `PostText` response. You can set the value between 5 and 86,400 seconds (24 hours).

Type: Integer

Valid Range: Minimum value of 5. Maximum value of 86400.

Required: No

turnsToLive

The number of conversation turns that the context should be active. A conversation turn is one `PostContent` or `PostText` request and the corresponding response from Amazon Lex.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 20.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Button

Service: Amazon Lex Runtime Service

Represents an option to be shown on the client platform (Facebook, Slack, etc.)

Contents

text

Text that is visible to the user on the button.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 15.

Required: Yes

value

The value sent to Amazon Lex when a user chooses the button. For example, consider button text "NYC." When the user chooses the button, the value sent can be "New York City."

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1000.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DialogAction

Service: Amazon Lex Runtime Service

Describes the next action that the bot should take in its interaction with the user and provides information about the context in which the action takes place. Use the `DialogAction` data type to set the interaction to a specific state, or to return the interaction to a previous state.

Contents

type

The next action that the bot should take in its interaction with the user. The possible values are:

- `ConfirmIntent` - The next action is asking the user if the intent is complete and ready to be fulfilled. This is a yes/no question such as "Place the order?"
- `Close` - Indicates that there will not be a response from the user. For example, the statement "Your order has been placed" does not require a response.
- `Delegate` - The next action is determined by Amazon Lex.
- `ElicitIntent` - The next action is to determine the intent that the user wants to fulfill.
- `ElicitSlot` - The next action is to elicit a slot value from the user.

Type: String

Valid Values: `ElicitIntent` | `ConfirmIntent` | `ElicitSlot` | `Close` | `Delegate`

Required: Yes

fulfillmentState

The fulfillment state of the intent. The possible values are:

- `Failed` - The Lambda function associated with the intent failed to fulfill the intent.
- `Fulfilled` - The intent has fulfilled by the Lambda function associated with the intent.
- `ReadyForFulfillment` - All of the information necessary for the intent is present and the intent ready to be fulfilled by the client application.

Type: String

Valid Values: `Fulfilled` | `Failed` | `ReadyForFulfillment`

Required: No

intentName

The name of the intent.

Type: String

Required: No

message

The message that should be shown to the user. If you don't specify a message, Amazon Lex will use the message configured for the intent.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

messageFormat

- `PlainText` - The message contains plain UTF-8 text.
- `CustomPayload` - The message is a custom format for the client.
- `SSML` - The message contains text formatted for voice output.
- `Composite` - The message contains an escaped JSON object containing one or more messages. For more information, see [Message Groups](#).

Type: String

Valid Values: `PlainText` | `CustomPayload` | `SSML` | `Composite`

Required: No

slots

Map of the slots that have been gathered and their values.

Type: String to string map

Required: No

slotToElicit

The name of the slot that should be elicited from the user.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

GenericAttachment

Service: Amazon Lex Runtime Service

Represents an option rendered to the user when a prompt is shown. It could be an image, a button, a link, or text.

Contents

attachmentLinkUrl

The URL of an attachment to the response card.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

buttons

The list of options to show to the user.

Type: Array of [Button](#) objects

Array Members: Minimum number of 0 items. Maximum number of 5 items.

Required: No

imageUrl

The URL of an image that is displayed to the user.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

subTitle

The subtitle shown below the title.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 80.

Required: No

title

The title of the option.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 80.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IntentConfidence

Service: Amazon Lex Runtime Service

Provides a score that indicates the confidence that Amazon Lex has that an intent is the one that satisfies the user's intent.

Contents

score

A score that indicates how confident Amazon Lex is that an intent satisfies the user's intent. Ranges between 0.00 and 1.00. Higher scores indicate higher confidence.

Type: Double

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IntentSummary

Service: Amazon Lex Runtime Service

Provides information about the state of an intent. You can use this information to get the current state of an intent so that you can process the intent, or so that you can return the intent to its previous state.

Contents

dialogActionType

The next action that the bot should take in its interaction with the user. The possible values are:

- `ConfirmIntent` - The next action is asking the user if the intent is complete and ready to be fulfilled. This is a yes/no question such as "Place the order?"
- `Close` - Indicates that there will not be a response from the user. For example, the statement "Your order has been placed" does not require a response.
- `ElicitIntent` - The next action is to determine the intent that the user wants to fulfill.
- `ElicitSlot` - The next action is to elicit a slot value from the user.

Type: String

Valid Values: `ElicitIntent` | `ConfirmIntent` | `ElicitSlot` | `Close` | `Delegate`

Required: Yes

checkpointLabel

A user-defined label that identifies a particular intent. You can use this label to return to a previous intent.

Use the `checkpointLabelFilter` parameter of the `GetSessionRequest` operation to filter the intents returned by the operation to those with only the specified label.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[a-zA-Z0-9-]+`

Required: No

confirmationStatus

The status of the intent after the user responds to the confirmation prompt. If the user confirms the intent, Amazon Lex sets this field to `Confirmed`. If the user denies the intent, Amazon Lex sets this value to `Denied`. The possible values are:

- `Confirmed` - The user has responded "Yes" to the confirmation prompt, confirming that the intent is complete and that it is ready to be fulfilled.
- `Denied` - The user has responded "No" to the confirmation prompt.
- `None` - The user has never been prompted for confirmation; or, the user was prompted but did not confirm or deny the prompt.

Type: String

Valid Values: `None` | `Confirmed` | `Denied`

Required: No

fulfillmentState

The fulfillment state of the intent. The possible values are:

- `Failed` - The Lambda function associated with the intent failed to fulfill the intent.
- `Fulfilled` - The intent has fulfilled by the Lambda function associated with the intent.
- `ReadyForFulfillment` - All of the information necessary for the intent is present and the intent ready to be fulfilled by the client application.

Type: String

Valid Values: `Fulfilled` | `Failed` | `ReadyForFulfillment`

Required: No

intentName

The name of the intent.

Type: String

Required: No

slots

Map of the slots that have been gathered and their values.

Type: String to string map

Required: No

slotToElicit

The next slot to elicit from the user. If there is not slot to elicit, the field is blank.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PredictedIntent

Service: Amazon Lex Runtime Service

An intent that Amazon Lex suggests satisfies the user's intent. Includes the name of the intent, the confidence that Amazon Lex has that the user's intent is satisfied, and the slots defined for the intent.

Contents

intentName

The name of the intent that Amazon Lex suggests satisfies the user's intent.

Type: String

Required: No

nlIntentConfidence

Indicates how confident Amazon Lex is that an intent satisfies the user's intent.

Type: [IntentConfidence](#) object

Required: No

slots

The slot and slot values associated with the predicted intent.

Type: String to string map

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ResponseCard

Service: Amazon Lex Runtime Service

If you configure a response card when creating your bots, Amazon Lex substitutes the session attributes and slot values that are available, and then returns it. The response card can also come from a Lambda function (`dialogCodeHook` and `fulfillmentActivity` on an intent).

Contents

contentType

The content type of the response.

Type: String

Valid Values: `application/vnd.amazonaws.card.generic`

Required: No

genericAttachments

An array of attachment objects representing options.

Type: Array of [GenericAttachment](#) objects

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Required: No

version

The version of the response card format.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SentimentResponse

Service: Amazon Lex Runtime Service

The sentiment expressed in an utterance.

When the bot is configured to send utterances to Amazon Comprehend for sentiment analysis, this field structure contains the result of the analysis.

Contents

sentimentLabel

The inferred sentiment that Amazon Comprehend has the highest confidence in.

Type: String

Required: No

sentimentScore

The likelihood that the sentiment was correctly inferred.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Document History for Amazon Lex

- **Latest documentation update:** September 9, 2021

The following table describes important changes in each release of Amazon Lex. For notification about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
New feature	Amazon Lex now supports the Korean (ko-KR) locale. For more information, see Languages supported by Amazon Lex .	September 9, 2021
New feature	Amazon Lex now supports the English (Indian) locale. For more information, see Languages supported in Amazon Lex .	July 15, 2021
New feature	Amazon Lex now provides a tool to migrate a bot to the Amazon Lex V2 API. For more information, see Migrating a bot .	July 13, 2021
New feature	Amazon Lex now supports the Japanese (Japan) locale. For more information, see Languages supported by Amazon Lex .	April 1, 2021
New feature	Amazon Lex now supports the German (German) (de-DE) and Spanish (Latin American) (es-419) locales. For more	November 23, 2020

information, see [Languages supported by Amazon Lex](#).

[New feature](#)

Amazon Lex now supports using contexts to manage activating intents. For more information, see [Setting Intent Context](#).

November 19, 2020

[New feature](#)

Amazon Lex now supports the French (fr-FR), French Canadian (fr-CA), Italian (it-IT) and Spanish (es-ES) locales. For a complete list of supported locales, see [Languages supported by Amazon Lex](#).

November 11, 2020

[New feature](#)

Amazon Lex now supports the Spanish (US) (es-US) locale. For more information, see [Languages supported by Amazon Lex](#).

September 22, 2020

[New feature](#)

Amazon Lex now supports the English (British) (en-GB) locale. For more information, see [Languages supported by Amazon Lex](#).

September 15, 2020

[New feature](#)

Amazon Lex now supports the English (Australian) (en-AU) locale. For more information, see [Languages supported by Amazon Lex](#).

September 8, 2020

New feature	Amazon Lex now has 7 new built-in intents and 9 new built-in slot types. For more information, see Built-in Intents and Slot Types .	September 8, 2020
New example	Learn how to create an Amazon Lex bot that customer support agents can use to answer customer questions by searching for answers with Amazon Kendra. For more information, see Example: Call Center Agent Assistant .	August 10, 2020
New feature	Amazon Lex can now return up to four alternative intents based on confidence scores. For more information, see Using Confidence Scores .	August 6, 2020
Region expansion	Amazon Lex is now available in Asia Pacific (Tokyo) (ap-northeast-1).	June 30, 2020
New feature	Amazon Lex now supports searching Amazon Kendra indexes for answers to frequently asked questions . For more information, see AMAZON.KendraSearchIntent .	June 11, 2020

New feature	Amazon Lex now returns more information in conversation logs. For more information, see Viewing Text Logs in Amazon CloudWatch Logs .	June 9, 2020
Region expansion	Amazon Lex is now available in Asia Pacific (Singapore) (ap-southeast-1) , Europe (Frankfurt) (eu-central-1), and Europe (London) (eu-west-2).	April 23, 2020
New feature	Amazon Lex now supports tagging. You can use tagging to identify resources, allocate costs, and control access. For more information, see Tagging your Amazon Lex Resources .	March 12, 2020
New feature	Amazon Lex now supports regular expressions for the AMAZON.AlphaNumeric built-in slot type. For more information, see AMAZON.AlphaNumeric .	February 6, 2020
New feature	Amazon Lex can now log conversation information and obfuscate slot values in those logs. For more information, see Creating Conversation Logs and Slot Obfuscation .	December 19, 2019
Region expansion	Amazon Lex is now available in Asia Pacific (Sydney) (ap-southeast-2).	December 17, 2019

New feature	Amazon Lex is now HIPAA compliant. For more information, see Compliance Validation for Amazon Lex .	December 10, 2019
New feature	Amazon Lex can now send user utterances to Amazon Comprehend to analyze the sentiment of the utterance . For more information, see Sentiment Analysis .	November 21, 2019
New feature	Amazon Lex is now SOC compliant. For more information, see Compliance Validation for Amazon Lex .	November 19, 2019
New feature	Amazon Lex is now PCI compliant. For more information, see Compliance Validation for Amazon Lex .	October 17, 2019
New feature	Added support for adding a checkpoint to an intent so that you can easily return to the intent during a conversation. For more information, see Managing Sessions .	October 10, 2019
New feature	Added support for the <code>AMAZON.FallbackIntent</code> so that your bot can handle situations when user input is not as expected. For more information, see AMAZON.FallbackIntent .	October 3, 2019

New feature	Amazon Lex enables you to manage session information for your bots. For more information, see Managing Sessions With the Amazon Lex API .	August 8, 2019
Region expansion	Amazon Lex is now available in US West (Oregon) (us-west-2).	May 8, 2018
New feature	Added support for exporting and importing in Amazon Lex format. For more information, see Importing and Exporting Amazon Lex Bots, Intents, and Slot Types .	February 13, 2018
New feature	Amazon Lex now supports additional response messages for bots. For more information, see Responses .	February 8, 2018
Region expansion	Amazon Lex is now available in Europe (Ireland) (eu-west-1).	November 21, 2017
New feature	Added support for deploying Amazon Lex bots on Kik. For more information, see Integrating an Amazon Lex Bot with Kik .	November 20, 2017
New feature	Added support for new built-in slot types and request attributes. For more information, see Built-in Slot Types and Setting Request Attributes .	November 3, 2017

New feature	Added export to Alexa Skills Kit feature. For more information, see Exporting to an Alexa Skill .	September 7, 2017
New feature	Added synonym support for slot type values. For more information, see Custom Slot Types .	August 31, 2017
New feature	Added AWS CloudTrail integration. For more information, see Monitoring Amazon Lex API Calls with AWS CloudTrail Logs .	August 15, 2017
Expanded documentation	Added Getting Started examples for the AWS CLI. For more information, see https://docs.aws.amazon.com/lex/latest/dg/gs-cli.html	May 22, 2017
New guide	This is the first release of the <i>Amazon Lex User Guide</i> .	April 19, 2017

AWS Glossary

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