
Amazon Transcribe

Developer Guide



Amazon Transcribe: Developer Guide

Copyright © 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

What is Amazon Transcribe?	1
Amazon Transcribe use cases	1
Use cases in action	2
Pricing	2
Supported languages	3
Character sets	5
Afrikaans	6
Arabic	7
Chinese, Simplified	9
Chinese, Traditional	10
Danish	10
Dutch	11
English	13
Farsi	14
French	15
German	17
Hebrew	18
Hindi	20
Indonesian	22
Italian	23
Japanese	24
Korean	25
Malay	26
Portuguese	27
Russian	29
Spanish	30
Tamil	31
Telugu	33
Thai	36
Turkish	38
Feature summary	40
Guidelines and quotas	42
Supported AWS Regions	42
Guidelines	42
Quotas	42
How it works	45
Data input and output	45
Media formats	45
Sample rates	46
Output	46
Transcribing numbers	48
Getting started	52
Signing up for an AWS account	52
Installing the AWS CLI and SDKs	52
Creating an IAM user	53
Creating an Amazon S3 bucket	53
Creating an IAM policy	53
Transcribing with the AWS Management Console	55
Transcribing with the AWS CLI	63
Starting a new transcription job	64
Getting the status of a transcription job	65
Listing your transcription jobs	65
Deleting your transcription job	66
Transcribing with the AWS SDKs	66

Working with AWS SDKs	74
Transcribing with HTTP or WebSockets	74
Streaming transcriptions	76
Streaming and partial results	76
Partial-result stabilization	77
Partial-result stabilization example output	79
Setting up a streaming transcription	80
Event stream encoding	80
Data frames	81
Setting up an HTTP/2 stream	82
Setting up a WebSocket stream	86
Job queueing	92
Enabling job queueing	92
Tagging resources	96
Tag-based access control	96
Adding tags to your Amazon Transcribe resources	97
Improving transcription accuracy	100
Custom vocabularies	101
Custom vocabulary tables versus lists	101
Creating a custom vocabulary using a table	102
Creating a custom vocabulary using a list	106
Using a custom vocabulary	110
Custom language models	115
Data sources	116
Training versus tuning data	116
Creating a custom language model	117
Using a custom language model	121
Alternative transcriptions	126
Requesting alternative transcriptions	127
Filtering words	131
Creating a vocabulary filter	132
Creating custom vocabulary filters	132
Using a custom vocabulary filter	135
Using a custom vocabulary filter in a batch transcription	110
Using a custom vocabulary filter in a streaming transcription	114
Redacting transcripts	142
Redacting PII in your batch job	143
Redacting or identifying PII in a real-time stream	146
Example PII redaction and identification output	149
Example redacted batch output	149
Example redacted streaming output	151
Example PII identification output	152
Transcribing multi-channel audio	154
Transcribing multi-channel audio	154
Transcribing multi-channel audio streams	157
Transcribing multi-channel audio in an HTTP/2 stream	157
Transcribing multi-channel audio in a WebSocket stream	158
Multi-channel streaming output	159
Identifying languages	160
Batch language identification	160
Identifying languages in multi-language audio	160
Improving language identification accuracy	161
Combining language identification with other Amazon Transcribe features	162
Using language identification with batch transcriptions	162
Streaming language identification	167
Using language identification with streaming media	168
Identifying speakers (diarization)	172

Identifying speakers in audio	172
Identifying speakers in real-time streams	173
HTTP/2 streaming	173
WebSocket streaming	174
Creating subtitles	175
Generating subtitle files	175
Analyzing call center audio	178
Call Analytics use cases	178
Considerations and additional information	179
Call Analytics insights	179
Call categorization	179
Call characteristics	180
Call summarization	180
Sensitive data redaction	181
Sentiment analysis	181
Creating categories	181
Rule criteria	187
Starting a Call Analytics job	189
Example output	194
Call categorization	194
Call characteristics	194
Call summarization	195
Sensitive data redaction	196
Sentiment analysis	196
Compiled output	198
Code examples	202
Actions	203
Create a custom vocabulary	203
Delete a custom vocabulary	204
Delete a medical transcription job	205
Delete a transcription job	206
Get a custom vocabulary	207
Get a transcription job	208
List custom vocabularies	208
List medical transcription jobs	209
List transcription jobs	210
Start a medical transcription job	212
Start a transcription job	213
Update a custom vocabulary	215
Scenarios	215
Create and refine a custom vocabulary	216
Transcribe audio and get job data	221
Cross-service examples	222
Build an Amazon Transcribe app	223
Build an Amazon Transcribe streaming app	223
Convert text to speech and back to text	224
Security	225
Identity and Access Management	225
Audience	226
Authenticating with identities	226
Managing access using policies	228
How Amazon Transcribe works with IAM	229
Confused deputy prevention	234
Identity-based policy examples	235
Troubleshooting	240
Data protection	242
Inter-network traffic privacy	242

Data encryption	243
Opting out of using your data for service improvement	245
Monitoring Amazon Transcribe	245
Monitoring with CloudWatch	246
Monitoring Amazon Transcribe with CloudTrail	246
Using Amazon EventBridge with Amazon Transcribe	249
Compliance Validation	252
Resilience	252
Infrastructure Security	253
Vulnerability analysis and management	253
VPC endpoints (AWS PrivateLink)	253
Security Best Practices	255
Amazon Transcribe Medical	256
How it works	256
Speech input	257
Streaming transcription overview	258
Batch transcription overview	258
Transcribing numbers	259
Transcribing medical terms and measurements	260
Getting started	262
Set up an AWS account	262
Getting started with the AWS Management Console (streaming)	263
Getting started with batch transcription	264
Streaming transcription	264
Event stream encoding	269
Using HTTP/2 streaming	270
Using WebSocket streaming	283
Transcribing a medical conversation	290
Transcribing an audio file	291
Transcribing a real-time stream	294
Identifying speakers	296
Transcribing multi-channel audio	302
Transcribing a medical dictation	307
Transcribing an audio file	308
Transcribing a streaming medical dictation	311
Creating and using medical custom vocabularies	313
Creating a text file for your medical custom vocabulary	313
Using a text file to create a medical custom vocabulary	316
Transcribing an audio file using a medical custom vocabulary	318
Transcribing a real-time stream using a medical custom vocabulary	319
Character sets for Amazon Transcribe Medical	321
Identifying PHI in a transcript	322
Identifying PHI in an audio file	323
Identifying PHI in a real-time stream	326
Generating alternative transcriptions	327
Guidelines and quotas	329
Supported AWS Regions	42
Guidelines	329
Quotas	329
VPC endpoints (AWS PrivateLink)	331
Considerations for Amazon Transcribe Medical VPC endpoints	331
Creating an interface VPC endpoint for Amazon Transcribe Medical	331
Creating a VPC endpoint policy for Amazon Transcribe Medical streaming	332
Document history	333
AWS glossary	340

What is Amazon Transcribe?

Amazon Transcribe is an automatic speech recognition service that uses machine learning models to convert audio to text. You can use Amazon Transcribe as a standalone transcription service or to add speech-to-text capabilities to any application.

With Amazon Transcribe, you can improve accuracy for your specific use case with language customization, filter content to ensure customer privacy or audience-appropriate language, analyze content in multi-channel audio, differentiate the speech of individual speakers, and more.

Important

Amazon Transcribe is covered under AWS's HIPAA eligibility and BAA which requires BAA customers to encrypt all PHI at rest and in transit when in use. Automatic PHI identification is available at no additional charge and in all regions where Amazon Transcribe operates. For more information, refer to [HIPAA eligibility and BAA](#).

To view a complete list of features, see [Amazon Transcribe features \(p. 40\)](#).

You can transcribe streaming media in real time or you can upload and transcribe media files. To see which languages are supported for each type of transcription, refer to the [Supported languages and language-specific features \(p. 3\)](#) table.

See [What is Amazon Transcribe?](#) for a short video tour of this service.

To learn more, see [How Amazon Transcribe works \(p. 45\)](#) and [Getting started with Amazon Transcribe \(p. 52\)](#).

Tip

Information on the **Amazon Transcribe API** is located in the [API Reference](#).

Amazon Transcribe use cases

Amazon Transcribe is a robust speech-to-text service that offers a diverse array of features, many of which can be combined between Amazon Transcribe and other AWS services.

- Gain insight into agent-customer calls using [Call Analytics \(p. 178\)](#). This feature automatically analyzes 11 different criteria without any customization on your part. For each speaker, you get sentiment data, talk time, non-talk time, loudness, interruptions, and talk speed. Call summarization, call categorization, and turn-by-turn output are provided for the whole call.

We've also launched two new analytics options for call center audio: [post-call analytics](#) (designed for audio files located in an Amazon S3 bucket) and [live-call analytics](#) (designed for live audio streams).

- Get a summary of customer-agent interactions with [call summarization \(p. 180\)](#), which provides an at-a-glance summary of issues, action items, and outcomes for every call.
- Teach Amazon Transcribe industry-specific terms, unique spelling, acronyms, and any words that are not being rendered correctly in your transcription results using [custom vocabularies \(p. 101\)](#). Providing Amazon Transcribe with custom vocabularies can improve the accuracy of your transcription output. See also: [Custom language models \(p. 115\)](#).
- Create [subtitles \(p. 175\)](#) for your video files. You can also use [content redaction \(p. 142\)](#) (only in US English) and [vocabulary filtering \(p. 131\)](#) when generating subtitles to ensure your content is audience-appropriate. Note that filtered or redacted content shows as white space, ***, or [P I I] in your transcript and subtitle files, but **the audio itself is not altered**.

- Redact personally identifiable information (PII), such as social security numbers, from your transcripts using standard [content redaction \(p. 142\)](#) or Call Analytics [sensitive data redaction \(p. 181\)](#). Call Analytics can also **redact your audio** by replacing spoken PII with silence.
- Identify individual speakers in an audio clip using [speaker diarization \(p. 172\)](#). When you activate speaker diarization, Amazon Transcribe attaches a unique attribute to each speaker in your transcription output.
- Remove proprietary terms from your transcript using [vocabulary filtering \(p. 131\)](#). For example, you can mask the name of a new product in a pre-launch stakeholder meeting. Vocabulary filtering can also be used to mask profane, offensive, or audience-inappropriate terms.
- Using multi-channel audio, you can have Amazon Transcribe produce a separate transcript for each channel, or have all channels transcribed in one output file. See [Transcribing multi-channel audio \(p. 154\)](#).
- If your audio is not in a language you speak, let Amazon Transcribe identify the language for you using [language identification \(p. 160\)](#). You can then use [Amazon Translate](#) to translate your transcript, and have [Amazon Polly](#) read your transcript back to you.
- Improve [streaming transcription accuracy](#) with [partial result stabilization \(p. 77\)](#), which can also be used to [adjust the latency of your transcript](#).

Tip

For use case code examples, refer to the [AWS Samples repository](#) on GitHub.

Use cases in action

Here are some diverse examples of how individuals and organizations are using Amazon Transcribe.

- [Live transcriptions of F1 races using Amazon Transcribe](#)
- [Generate high-quality meeting notes using Amazon Transcribe and Amazon Comprehend](#)
- [Boost transcription accuracy of class lectures with custom language models for Amazon Transcribe](#)
- [Make your audio and video files searchable using Amazon Transcribe and Amazon Kendra](#)
- [Perform medical transcription analysis in real-time with AWS AI services and Twilio Media Streams](#)

Pricing

Amazon Transcribe is a pay-as-you-go service; pricing is based on seconds of transcribed audio, billed on a monthly basis. For more information on cost, including cost-breakdown examples for various AWS Regions, see [Amazon Transcribe Pricing](#).

Supported languages and language-specific features

The languages supported by Amazon Transcribe are listed in the following table; also listed are the features that are language-specific. Please verify that the feature you wish to use is supported for the language in your media before proceeding with your transcription.

To view the complete list of Amazon Transcribe features, refer to the [Feature summary \(p. 40\)](#).

In the following table, 'batch' refers to transcribing a media file located in an Amazon S3 bucket and 'streaming' refers to transcribing streamed media in real time.

Language	Language Code	Data input (p. 45)	Transcribing numbers (p. 46)	Acronyms (p. 47)	Custom language models (p. 177)	Redaction (p. 178)	Call Analytics (p. 178)
Afrikaans (p. 6)	af-ZA	batch	no	batch	no	no	no
Arabic (p. 7), Gulf	ar-AE	batch	no	no	no	no	batch
Arabic (p. 7), Modern Standard	ar-SA	batch	no	no	no	no	no
Chinese, Simplified (p. 9)	zh-CN	batch, streaming	no	no	no	no	batch
Chinese, Traditional (p. 10)	zh-TW	batch	no	no	no	no	no
Danish (p. 10)	da-DK	batch	no	batch	no	no	no
Dutch (p. 11)	nl-NL	batch	no	batch	no	no	no
English (p. 13), Australian	en-AU	batch, streaming	batch, streaming	batch, streaming	batch	no	batch
English (p. 13), British	en-GB	batch, streaming	batch, streaming	batch, streaming	batch, streaming	no	batch
English (p. 13), Indian	en-IN	batch	batch	batch	no	no	batch
English (p. 13), Irish	en-IE	batch	batch	batch	no	no	batch
English (p. 13), New Zealand	en-NZ	batch	batch	batch	no	no	no
English (p. 13), Scottish	en-AB	batch	batch	batch	no	no	batch

Language	Language Code	Data input (p. 45)	Transcribing numbers (p. 178)	Acronyms (p. 178)	Custom language models (p. 178)	Redaction (p. 178)	Call Analytics (p. 178)
English (p. 13) South African	en-ZA	batch	batch	batch	no	no	no
English (p. 13) US	en-US	batch, streaming	batch, streaming	batch, streaming	batch, streaming	batch, streaming	batch
English (p. 13) Welsh	en-WL	batch	batch	batch	no	no	batch
French (p. 15)	fr-FR	batch, streaming	no	batch, streaming	no	no	batch
French (p. 15) Canadian	fr-CA	batch, streaming	no	batch, streaming	no	no	batch
Farsi (p. 14)	fa-IR	batch	no	no	no	no	no
German (p. 17)	de-DE	batch, streaming	batch, streaming	batch, streaming	no	no	batch
German (p. 17) Swiss	de-CH	batch	batch	batch	no	no	batch
Hebrew (p. 18)	he-IL	batch	no	no	no	no	no
Hindi (p. 20) Indian	hi-IN	batch	no	batch	batch	no	batch
Indonesian (p. 22)	id-ID	batch	no	batch	no	no	no
Italian (p. 23)	it-IT	batch, streaming	no	batch, streaming	no	no	batch
Japanese (p. 24)	ja-JP	batch, streaming	no	no	no	no	batch
Korean (p. 25)	ko-KR	batch, streaming	no	no	no	no	batch
Malay (p. 26)	ms-MY	batch	no	batch	no	no	no
Portuguese (p. 27)	pt-BR	batch	no	batch	no	no	batch
Portuguese (p. 27) Brazilian	pt-BR	batch, streaming	no	batch, streaming	no	no	batch
Russian (p. 29)	ru-RU	batch	no	no	no	no	no
Spanish (p. 30)	es-ES	batch	no	batch	no	no	batch
Spanish (p. 30) US	es-US	batch, streaming	no	batch, streaming	batch	no	batch
Tamil (p. 31)	ta-IN	batch	no	no	no	no	no
Telugu (p. 33)	te-IN	batch	no	no	no	no	no

Language	Language Code	Data input (p. 45)	Transcribing numbers (p. 46)	Acronyms (p. 47)	Custom language models (p. 177)	Redaction (p. 178)	Call Analytics (p. 178)
Thai (p. 36)	th-TH	batch	no	batch	no	no	no
Turkish (p. 38)	tr-TR	batch	no	batch	no	no	no

Character sets for custom vocabularies and vocabulary filters

For each language Amazon Transcribe supports, there is a specific set of characters Amazon Transcribe can recognize. When you create a custom vocabulary or vocabulary filter, use only the characters listed in your language's character set. If you use unsupported characters, your custom vocabulary or vocabulary filter fails.

Important

Be sure to check that your custom vocabulary file uses only the supported Unicode code points and code point sequences listed within the following character sets.

Many Unicode characters can appear identical in popular fonts, even if they use different code points. **Only the code points listed in this guide are supported.** For example, the French word *déjà* can be rendered using *precomposed* characters (where one Unicode value represents an accented character) or *decomposed* characters (where two Unicode values represent an accented character, one value for the base character and another for the accent).

- **Precomposed version:** 0064 00E9 006A 00E0 (renders as *déjà*)
- **Decomposed version:** 0064 0065 0301 006A 0061 0300 (renders as *déjà*)

Topics

- [Afrikaans character set \(p. 6\)](#)
- [Arabic character set \(p. 7\)](#)
- [Chinese, Mandarin \(Mainland China\), Simplified character set \(p. 9\)](#)
- [Chinese, Mandarin \(Taiwan\), Traditional character set \(p. 10\)](#)
- [Danish character set \(p. 10\)](#)
- [Dutch character set \(p. 11\)](#)
- [English character set \(p. 13\)](#)
- [Farsi character set \(p. 14\)](#)
- [French character set \(p. 15\)](#)
- [German character set \(p. 17\)](#)
- [Hebrew character set \(p. 18\)](#)
- [Hindi character set \(p. 20\)](#)
- [Indonesian character set \(p. 22\)](#)
- [Italian character set \(p. 23\)](#)
- [Japanese character set \(p. 24\)](#)
- [Korean character set \(p. 25\)](#)
- [Malay character set \(p. 26\)](#)
- [Portuguese character set \(p. 27\)](#)

- [Russian character set \(p. 29\)](#)
- [Spanish character set \(p. 30\)](#)
- [Tamil character set \(p. 31\)](#)
- [Telugu character set \(p. 33\)](#)
- [Thai character set \(p. 36\)](#)
- [Turkish character set \(p. 38\)](#)

Afrikaans character set

For Afrikaans custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
á	00E1	ï	00EF
è	00E8	ó	00F3
é	00E9	ô	00F4
ê	00EA	ö	00F6
ë	00EB	ú	00FA
í	00ED	û	00FB
î	00EE	ü	00FC

You can use the following International Phonetic Alphabet characters in the `IPA` field of the custom vocabulary input file:

Character	Code	Character	Code
a	0061	z	007A
b	0062	æ	00E6
c	0063	ð	00F0
d	0064	ø	00F8
e	0065	ɿ	014B
e~	0065 0303	œ	0153
f	0066	ɐ	0250
g	0067	ɑ	0251

Character	Code	Character	Code
h	0068	ħ	0251 0303
i	0069	ɨ	0252
j	006A	ɟ	0254
k	006B	ķ	0254 0303
l	006C	ə	0259
m	006D	ɛ	025B
n	006E	ĕ	025B 0303
o	006F	ɜ	025C
o~	006F 0303	ɪ	026A
p	0070	†	026C
r	0072	ɹ	0279
s	0073	ʁ	0281
t	0074	ʃ	0283
u	0075	ʊ	028A
v	0076	ʌ	028C
w	0077	ɜ	0292
x	0078	θ	03B8
y	0079		

Arabic character set

For Arabic custom vocabularies, you can use the following Unicode characters in the `Phrase` and `SoundsLike` fields. You can also use the hyphen (-) character to separate words.

Character	Code	Character	Code
ء	0621	س	0633
آ	0622	ش	0634
أ	0623	ص	0635
ؤ	0624	ض	0636
إ	0625	ط	0637
ئ	0626	ظ	0638
ا	0627	ع	0639
ب	0628	غ	063A

Character	Code	Character	Code
ة	0629	ف	0641
ت	062A	ق	0642
ث	062B	ك	0643
ج	062C	ل	0644
ح	062D	م	0645
خ	062E	ن	0646
د	062F	ه	0647
ذ	0630	و	0648
ر	0631	ى	0649
ز	0632	ي	064A

You can use the following International Phonetic Alphabet characters in the IPA field of the custom vocabulary input file:

Character	Code	Character	Code
a	0061	tʃ	0074 02E4
a:	0061 02D0	u	0075
b	0062	u:	0075 02D0
d	0064	v	0076
dʃ	0064 02E4	w	0077
f	0066	x	0078
h	0068	z	007A
i	0069	zʃ	007A 02E4
i:	0069 02D0	ð	00F0
j	006A	ðʃ	00F0 02E4
k	006B	ħ	0127
l	006C	ɣ	0263
m	006D	ɪ	026A
n	006E	ɰ	026B
p	0070	ʃ	0283
q	0071	ʒ	0292
r	0072	ʔ	0294

Character	Code	Character	Code
s	0073	ʃ	0295
sʃ	0073 02E4	θ	03B8
t	0074	χ	03C7

Chinese, Mandarin (Mainland China), Simplified character set

For Chinese (Simplified) custom vocabularies, the `Phrase` field can use any of the characters listed in the following file on GitHub.

- [zh-cn-character-set.txt](#)

The `SoundsLike` field can contain the pinyin syllables listed in the following file on GitHub.

- [pinyin-set.txt](#)

When you use pinyin syllables in the `SoundsLike` field, separate the syllables with a hyphen (-).

Amazon Transcribe represents the four tones in Chinese (Simplified) using numbers. The following table shows how tone marks are mapped for the word 'ma'.

Tone	Tone mark	Tone number
Tone 1	mā	ma1
Tone 2	má	ma2
Tone 3	mǎ	ma3
Tone 4	mà	ma4

Note

For the 5th (neutral) tone, you can use Tone 1, with the exception of 'er', which must be mapped to Tone 2. For example, 打转儿 would be represented as 'da3-zhuan4-er2'.

Chinese (Simplified) custom vocabularies don't use the `IPA` field, but you must still include the `IPA` header in the custom vocabulary table.

The following example is an input file in text format. The example uses spaces to align the columns. Your input files should use TAB characters to separate the columns. Include spaces only in the `DisplayAs` column.

Phrase	SoundsLike	IPA	DisplayAs
##	kang1-jian4		
##	qian3-ze2		
####	guo2-fang2-da4-chen2		
#####	shi4-jie4-bo2-lan3-hui4	###	

Chinese, Mandarin (Taiwan), Traditional character set

For Chinese (Traditional) custom vocabularies, the `Phrase` field can use any of the characters listed in the following file on GitHub.

- [zh-tw-character-set.txt](#)

The `SoundsLike` field can contain the zhuyin syllables listed in the following file on GitHub.

- [zhuyin-set.txt](#)

When you use zhuyin syllables in the `SoundsLike` field, separate the syllables with a hyphen (-).

Amazon Transcribe represents the four tones in Chinese (Traditional) using numbers. The following table shows how tone marks are mapped for the word ㄇㄩˊ.

Tone	Tone mark	
Tone 1	ㄇㄩ	
Tone 2	ㄇㄩˊ	
Tone 3	ㄇㄩˇ	
Tone 4	ㄇㄩˋ	

Chinese (Traditional) custom vocabularies don't use the `IPA` field, but you must still include the `IPA` header in the custom vocabulary table.

The following example is an input file in text format. The example uses spaces to align the columns. Your input files should use TAB characters to separate the columns. Include spaces only in the `DisplayAs` column.

Phrase	SoundsLike	IPA	DisplayAs
##	###-##		
##	###~-###		
####	###-###-###-##~		
#####	##-#####-###-##~-#####	###	

Danish character set

For Danish custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
Å	00C5	æ	00E6
Æ	00C6	é	00E9
Ø	00D8	ø	00F8
å	00E5		

You can use the following International Phonetic Alphabet characters in the IPA field of the custom vocabulary input file:

Character	Code	Character	Code
a	0061	æ	00E6
b	0062	ð	00F0
d	0064	ø	00F8
e	0065	ɲ	014B
f	0066	œ	0153
h	0068	ɐ	0250
i	0069	ɑ	0251
j	006A	ɒ	0252
k ^h	006B 02B0	ɔ	0254
l	006C	ɸ	0255
m	006D	ə	0259
n	006E	ɛ	025B
o	006F	g	0261
p ^h	0070 02B0	ɪ	026A
s	0073	ɸ	0281
t ^s	0074 02E2	ʊ	028A
u	0075	ʊ	028B
w	0077	ʌ	028C
y	0079		

Dutch character set

For Dutch custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
à	00E0	î	00EE
á	00E1	ï	00EF
â	00E2	ñ	00F1
ä	00E4	ò	00F2
ç	00E7	ó	00F3
è	00E8	ô	00F4
é	00E9	ö	00F6
ê	00EA	ù	00F9
ë	00EB	ú	00FA
ì	00EC	û	00FB
í	00ED	ü	00FC

You can use the following International Phonetic Alphabet characters in the `IPA` field of the custom vocabulary input file:

Character	Code	Character	Code
a:	0061 003A	ø:	00F8 003A
b:	0062 02D0	ŋ	014B
b	0062	œy	0153 0079
d	0064	œ:	0153 02D0
e:	0065 02D0	ɑ	0251
f	0066	ɔ	0254
g	0067	ɔu	0254 0075
i	0069	ɔ:	0254 02D0
j	006A	ə	0259
k	006B	ɛ	025B
l	006C	ɛ:	025B 003A

Character	Code	Character	Code
m	006D	ɛi	025B 0069
n	006E	ɦ	0266
o:	006F 02D0	ɪ	026A
p	0070	ɲ	0272
s	0073	ɾ	027E
t	0074	ʃ	0283
u	0075	ɣ	028F
v	0076	ʒ	0292
w	0077	χ	03C7
y	0079	ʏ	0263
z	007A		

English character set

For English custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can use the following International Phonetic Alphabet characters in the `IPA` field of the custom vocabulary input file:

Character	Code	Character	Code
au	0061 028A	w	0077
aɪ	0061 026A	z	007A
b	0062	æ	00E6
d	0064	ð	00F0
eɪ	0065 026A	ŋ	014B
f	0066	ɑ	0251
g	0067	ɔ	0254
h	0068	ɔɪ	0254 026A
i	0069	ə	0259

Character	Code	Character	Code
j	006A	ε	025B
k	006B	Ʒ	025D
l	006C	g	0261
l	006C 0329	ɪ	026A
m	006D	ɹ	0279
n	006E	ʃ	0283
ŋ	006E 0329	ʊ	028A
oʊ	006F 028A	ʌ	028C
p	0070	ɱ	028D
s	0073	ʒ	0292
t	0074	ɟʒ	02A4
u	0075	ʈʂ	02A7
v	0076	θ	03B8

Farsi character set

For Farsi custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields.

Character	Code	Character	Code
ء	0621	ظ	0638
آ	0622	ع	0639
أ	0623	غ	063A
ؤ	0624	ف	0641
ئ	0626	ق	0642
ا	0627	ج	0644
ب	0628	م	0645
ت	062A	ن	0646
ث	062B	ه	0647
ج	062C	و	0648
ح	062D	ـ	064E
خ	062E	ء	064F
د	062F	ـ	0650

Character	Code	Character	Code
ذ	0630	ّ	0651
ر	0631	پ	067E
ز	0632	چ	0686
س	0633	ژ	0698
ش	0634	ک	06A9
ص	0635	گ	06AF
ض	0636	ی	06CC
ط	0637		

You can use the following International Phonetic Alphabet in the IPA field of your custom vocabulary file:

Character	Code	Character	Code
b	0062	u	0075
d	0064	v	0076
f	0066	z	007A
g	0067	æ	00E6
h	0068	ɒ	0252
i	0069	ɛ	025B
j	006A	r	027E
k	006B	ʁ	0281
l	006C	ʃ	0283
m	006D	ʒ	0292
n	006E	ʔ	0294
o	006F	ʔ	0294
p	0070	ɔ̃	02A4
s	0073	tʃ	02A7
t	0074	χ	03C7

French character set

For French custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z

- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
À	00C0	à	00E0
Â	00C2	â	00E2
Ç	00C7	ç	00E7
È	00C8	è	00E8
É	00C9	é	00E9
Ê	00CA	ê	00EA
Ë	00CB	ë	00EB
Î	00CE	î	00EE
Ï	00CF	ï	00EF
Ô	00D4	ô	00F4
Ö	00D6	ö	00F6
Ù	00D9	ù	00F9
Û	00DB	û	00FB
Ü	00DC	ü	00FC

You can use the following International Phonetic Alphabet in the `IPA` field of your custom vocabulary file:

Character	Code	Character	Code
a	0061	z	007A
b	0062	ã	00E3
d	0064	ö	00F5
e	0065	ø	00F8
f	0066	ŋ	014B
i	0069	œ	0153
j	006A	œ̃	0153 0303
k	006B	e	0250

Character	Code	Character	Code
l	006C	ɔ	0254
m	006D	ə	0259
n	006E	ɛ	025B
o	006F	g	0261
p	0070	ɥ	0265
s	0073	ɲ	0272
t	0074	ʁ	0281
u	0075	ʃ	0283
v	0076	ʒ	0292
w	0077	ě	1EBD
y	0079		

German character set

For German custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
ä	00E4	Ä	00C4
ö	00F6	Ö	00D6
ü	00FC	Ü	00DC
ß	00DF		

You can use the following International Phonetic Alphabet characters in the `IPA` field of the custom vocabulary input file:

Character	Code	Character	Code
a	0061	ts	0074 0073

Character	Code	Character	Code
א	0061 026A	ו:	0075 02D0
אׁ	0061 028A	ׁ	0076
א:	0061 02D0	כ	0078
ב	0062	כׁ	007A
ד	0064	כ׃	0079 02D0
ע:	0065 02D0	כּ	00E3
פ	0066	כּׁ	00E7
ג	0067	כּ׃	00F8 02D0
ה	0068	כּׁׁ	014B
י:	0069 02D0	כּׁ׃	0153
יׁ	006A	כּׁׁׁ	0250 032F
כ	006B	כּׁׁ׃	0254
ל	006C	כּׁׁׁׁ	0254 028F
יׁׁ	006C 0329	כּׁׁׁ׃	0259
מ	006D	כּׁׁׁׁׁ	025B
מׁ	006D 0329	כּׁׁׁׁ׃	025B 02D0
נ	006E	כּׁׁׁׁׁׁ	026A
נׁ	006E 0329	כּׁׁׁׁׁ׃	0281
ו:	006F 02D0	כּׁׁׁׁׁׁׁ	0283
פ	0070	כּׁׁׁׁׁׁׁׁ	028A
פׁ	0070 0066	כּׁׁׁׁׁׁׁ׃	028F
ס	0073	כּׁׁׁׁׁׁׁׁׁ	02A7
ט	0074		

Hebrew character set

For Hebrew custom vocabularies, you can use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
-	002D	◊	05DD
כ	05D0	מ	05DE
כּ	05D1	ל	05DF

Character	Code	Character	Code
א	05D2	ג	05E0
ב	05D3	ד	05E1
ה	05D4	ו	05E2
ו	05D5	ז	05E3
ז	05D6	ח	05E4
ח	05D7	ט	05E5
ט	05D8	י	05E6
י	05D9	כ	05E7
ך	05DA	ל	05E8
כ	05DB	מ	05E9
ל	05DC	נ	05EA

You can use the following International Phonetic Alphabet characters in the IPA field of the custom vocabulary input file:

Character	Code	Character	Code
a	0061	p	0070
b	0062	s	0073
d	0064	t	0074
e	0065	u	0075
f	0066	v	0076
g	0067	w	0077
h	0068	z	007A
i	0069	ŋ	014B
j	006A	ʒ	0263
k	006B	ʃ	0283
l	006C	ʒ	0292
m	006D	ʔ	0294
n	006E	χ	03C7
o	006F		

Hindi character set

For Hindi custom vocabularies, you can use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
-	002D	थ	0925
.	002E	द	0926
ु	0901	ध	0927
.	0902	न	0928
:	0903	प	092A
अ	0905	फ	092B
आ	0906	ब	092C
इ	0907	भ	092D
ई	0908	म	092E
उ	0909	य	092F
ऊ	090A	र	0930
ऋ	090B	ल	0932
ए	090F	व	0935
ऐ	0910	श	0936
ऑ	0911	ष	0937
ओ	0913	स	0938
औ	0914	ह	0939
क	0915	ट	093E
ख	0916	ि	093F
ग	0917	ी	0940
घ	0918	ु	0941
ङ	0919	ू	0942
च	091A	े	0943
छ	091B	े	0945
ज	091C	ै	0947
झ	091D	ै	0948
ञ	091E	ँ	0949
ट	091F	ो	094B

Character	Code	Character	Code
ठ	0920	ै	094C
ड	0921	ँ	094D
ढ	0922	ज़	095B
ण	0923	झ	095C
त	0924	ढ	095D

Amazon Transcribe maps the following characters:

Character	Mapped to
न (0929)	न (0928)
र (0931)	र (0930)
क (0958)	क (0915)
ख (0959)	ख (0916)
ग (095A)	ग (0917)
फ (095E)	फ (092B)
य (095F)	य (092F)

You can use the following International Phonetic Alphabet characters in the IPA field of your input file:

Character	Code	Character	Code
a:	0097 0720	ŋ	0331
b	0098	ɖ	0598
b ^h	0098 0689	ɔ:	0596 0720
d	0100	ɖ ^h	0598 0689
d ^h	0100 0689	ə	0601
e:	0101 0720	ɛ:	0603 0720
f	0102	g	0609
i:	0105 0720	g ^h	0609 0689
j	0106	ɦ	0614
k	0107	ɪ	0618
k ^h	0107 0688	ɲ	0626
l	0108	ɳ	0627
m	0109	ɾ	0638

Character	Code	Character	Code
n	0110	ɣ	0642
o:	0111 0720	ʃ	0643
p	0112	ʈ	0648
p ^h	0112 0688	ʈ ^h	0648 0688
r	0114	ʊ	0650
s	0115	ʋ	0651
t	0116	ɖʒ	0676
t ^h	0116 0688	ɖʒ ^h	0676 0689
u:	0117 0720	ʈʃ	0679
z	0122	ʈ ^h	0679 0688

Indonesian character set

For Indonesian custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can use the following International Phonetic Alphabet characters in the `IPA` field of your input file:

Character	Code	Character	Code
a	0061	r	0072
ai	0061 0069	s	0073
au	0061 0075	t	0074
b	0062	ʈʃ	0074 0283
d	0064	u	0075
ɖʒ	0064 0292	v	0076
e	0065	w	0077
f	0066	x	0078
h	0068	y	0079
i	0069	ŋ	014B

Character	Code	Character	Code
j	006A	ɔ	0254
k	006B	ə	0259
l	006C	ε	025B
m	006D	g	0261
n	006E	γ	0263
o	006F	ɪ	026A
oì	006F 0069 032F	ɲ	0272
p	0070	ʃ	0283
q	0071	ʊ	028A

Italian character set

For Italian custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
À	00C0	à	00E0
Ä	00C4	ä	00E4
Ç	00C7	ç	00E7
È	00C8	è	00E8
É	00C9	é	00E9
Ê	00CA	ê	00EA
Ë	00CB	ë	00EB
Ì	00CC	ì	00EC
Ò	00D2	ò	00F2
Ù	00D9	ù	00F9
Ü	00DC	ü	00FC

You can use the following International Phonetic Alphabet characters in the IPA field of your input file:

Character	Code	Character	Code
a	0061	ss	0073 0073
b	0062	t	0074
bb	0062 0062	tt	0074 0074
d	0064	u	0075
dd	0064 0064	v	0076
e	0065	vv	0076 0076
f	0066	w	0077
ff	0066 0066	z	007A
gg	0067 0067	ɔ	0254
i	0069	ε	025B
j	006A	g	0261
k	006B	ɲ	0272
kk	006B 006B	ɲɲ	0272 0272
l	006C	ʃ	0283
ll	006C 006C	ʃʃ	0283 0283
m	006D	ʌ	028E
mm	006D 006D	ʌʌ	028E 028E
n	006E	ɖz	02A3
nn	006E 006E	ɖzɖz	02A3 02A3
o	006F	ɖʒ	02A4
p	0070	ɖʒɖʒ	02A4 02A4
pp	0070 0070	ʈs	02A6
r	0072	ʈsʈs	02A6 02A6
rr	0072 0072	ʈʃ	02A7
s	0073	ʈʃʈʃ	02A7 02A7

Japanese character set

For Japanese custom vocabularies, the `Phrase` and `SoundsLike` fields can use any of the characters listed in the following file on GitHub.

- [ja-jp-character-set.txt](#)

You can use the following International Phonetic Alphabet characters in the IPA field of your input file:

Character	Code	Character	Code
a	0061	p	0070
a:	0061 02D0	s	0073
b	0062	t	0074
d	0064	ts	0074 0073
dz	0064 007A	tɕ	0074 0255
dʒ	0064 0291	w	0077
e	0065	z	007A
e:	0065 02D0	ɕ	00E7
g	0067	ŋ	014B
h	0068	ɕ	0255
i	0069	ɯ	026F
i:	0069 02D0	ɯ:	026F 02D0
j	006A	ɲ	0274
k	006B	ɸ	0278
m	006D	ɾ	027E
n	006E	ʐ	0291
o	006F	ʔ	0294
o:	006F 02D0		

Korean character set

For Korean custom vocabularies, you can use any of the Hangul syllables in the `Phrase` and `SoundsLike` fields. For more information, see [Hangul Syllables](#) on Wikipedia.

You can use the following International Phonetic Alphabet characters in the IPA field of your input file:

Character	Code	Character	Code
a	0061	s	0073
e	0065	s#	0073 0348
h	0068	t	0074
i	0069	tɕ	0074 0255
je	006A 0065	tɕ ^h	0074 0255 02B0
jo	006A 006F	t ^h	0074 02B0

Character	Code	Character	Code
ju	006A 0075	t#	0074 0348
jε	006A 025B	ʔ	0074 031A
jʌ	006A 028C	t#ç	0074 0348 0255
ja	006A 0061	u	0075
k	006B	we	0077 0065
k ^h	006B 02B0	wi	0077 0069
k#	006B 0348	wε	0077 025B
k̄	006B 031A	wʌ	0077 028C
l	006C	wa	0077 0061
m	006D	ø	00F8
n	006E	ɲ	0014B
o	006F	ε	0025B
p	0070	ω	026F
p ^h	0070 02B0	wi	006F 0069
p#	0070 0348	r	027E
p̄	0070 031A	ʌ	028C

Malay character set

For Malay custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can use the following International Phonetic Alphabet characters in the `IPA` field of your input file:

Character	Code	Character	Code
F	0046	r	0072
a	0061	s	0073
ai	0061 0069	t	0074
au	0061 0075	tʃ	0074 0283

Character	Code	Character	Code
b	0062	v	0076
d	0064	w	0077
dʒ	0064 0292	x	0078
e	0065	y	0079
h	0068	ɲ	014B
i	0069	ɔ	0254
j	006A	ə	0259
k	006B	ɛ	025B
l	006C	g	0261
m	006D	ɣ	0263
n	006E	ɪ	026A
o	006F	ɲ	0272
ɔj	006F 0069 32F	ʃ	0283
p	0070	ʊ	028A
q	0071	ʊi	028A 0069

Portuguese character set

For Portuguese custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
À	00C0	à	00E0
Á	00C1	á	00E1
Â	00C2	â	00E2
Ã	00C3	ã	00E3
Ä	00C4	ä	00E4

Character	Code	Character	Code
Ç	00C7	ç	00E7
È	00C8	è	00E8
É	00C9	é	00E9
Ê	00CA	ê	00EA
Ë	00CB	ë	00EB
Í	00CD	í	00ED
Ñ	00D1	ñ	00F1
Ó	00D3	ó	00F3
Ô	00D4	ô	00F4
Õ	00D5	õ	00F5
Ö	00D6	ö	00F6
Ú	00DA	ú	00FA
Ü	00DC	ü	00FC

You can use the following International Phonetic Alphabet characters in the IPA field of your input file:

Character	Code	Character	Code
a	0061	s	0073
b	0062	t	0074
d	0064	tʃ	0074 0283
e	0065	u	0075
ě	1EBD	ũ	00169
f	0066	v	0076
g	0067	w	0077
i	0069	w~	0077 0303
ĩ	00129	z	007A
j	006A	ž	0250 0303
ĵ	006A 0303	ɔ	0254
k	006B	ε	025B
l	006C	ʝ	0272
m	006D	r	027E
n	006E	ʀ	0281

Character	Code	Character	Code
o	006F	ŷ	0283
ö	00F5	ʌ	028E
p	0070	ʒ	0292
r	0072	ɖʒ	02A4

Russian character set

For Russian custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
'	0027	п	043F
-	002D	р	0440
.	002E	с	0441
а	0430	т	0442
б	0431	у	0443
в	0432	ф	0444
г	0433	х	0445
д	0434	ц	0446
е	0435	ч	0447
ж	0436	ш	0448
з	0437	щ	0449
и	0438	ъ	044A
й	0439	ы	044B
к	043A	ь	044C
л	043B	э	044D
м	043C	ю	044E
н	043D	я	044F
о	043E	ë	0451

You can use the following International Phonetic Alphabet characters in the `IPA` field of your input file:

Character	Code	Character	Code
b	0062	t	0074

Character	Code	Character	Code
b ^j	0062 02B2	tʃ	0074 0283
d	0064	ʎ	0074 02B2
d ^j	0064 02B2	u	0075
f	0066	v	0076
f ^j	0066 02B2	v ^j	0076 02B2
g	0067	x	0078
g ^j	067 02B2	x ^j	0078 02B2
i	0069	z	007A
j	006A	z ^j	007A 02B2
k	006B	æ	00E6
k ^j	006B 02B2	ə	0259
l	006C	ɛ	025B
l ^j	006C 02B2	ɨ	0268
m	006D	ʃ	0283
m ^j	006D 02B2	ʃ ^j	0283 02B2
n	006E	ʊ	028A
n ^j	006E 02B2	ʌ	028C
p	0070	ɜ	0292
p ^j	0070 02B2	'i	02C8 0069
r	0072	'o	02C8 006F
r ^j	0072 02B2	'v	02C8 0075
s	0073	'ɛ	02C8 025B
s ^j	0073 02B2	'i	02C8 0268
ts	0074 0073	'a	02C8 0061

Spanish character set

For Spanish custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
Á	00C1	á	00E1
É	00C9	é	00E9
Í	00CD	ë	00ED
Ó	00D3	ó	0XF3
Ú	00DA	ú	00FA
Ñ	00D1	ñ	0XF1
ü	00FC		

You can use the following International Phonetic Alphabet characters in the `IPA` field of your input file:

Character	Code	Character	Code
a	0061	r	0072
b	0062	s	0073
d	0064	t	0074
e	0065	u	0075
f	0066	v	0076
g	0067	w	0077
h	0068	x	0078
i	0069	z	007A
j	006A	ŋ	014B
k	006B	ɲ	0272
l	006C	ɾ	027E
m	006D	ʃ	0283
n	006E	ɟ	029D
o	006F	tʃ	02A7
p	0070	θ	03B8

Tamil character set

For Tamil custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
அ	OB85	ர	OBBO
ஆ	OB86	ல	OB22
இ	OB87	வ	OB25
ஈ	OB88	ழ	OB24
உ	OB89	ள	OB23
ஊ	OB8A	ற	OB21
எ	OB8E	ன	OB29
ஏ	OB8F	ஐ	OB2C
ஐ	OB90	#	OB26
ஓ	OB92	ஷ	OB27
ஔ	OB93	ஸ	OB28
ஒள	OB94	ஹ	OB29
ஃ	OB83	.	OB2D
க	OB95	ா	OB2E
ங	OB99	ி	OB2F
ச	OB9A	ு	OB20
ஞ	OB9E	ு	OB21
ட	OB9F	ு	OB22
ண	OB23	ெ	OB26
த	OB24	ே	OB27
ந	OB28	ை	OB28
ப	OB2A	ொ	OB2A
ம	OB2E	ோ	OB2B
ய	OB2F	ௌ	OB2C

You can use the following International Phonetic Alphabet characters in the IPA field of your input file:

Character	Code	Character	Code
a	0061	v	0076
a:	0061 02D0	w	0077
b	0062	z	007A
d	0064	æ	00E6

Character	Code	Character	Code
dʒ	0064 0292	ð	00F0
e	0065	ŋ	014B
f	0066	ɑ	0251
g	0067	ɔ	0254
h	0068	ə	0259
i	0069	ɛ	025B
i:	0069 02D0	g	0261
j	006A	ɪ	026A
k	006B	l̪	026D
l	006C	ɳ	0272
m	006D	ɳ	0273
n	006E	ɻ	0279
ɳ	006E 032A	ɻ	0279
o	006F	ɻ	0279 0329
o:	006F 02D0	ɾ	027E
p	0070	ʂ	0282
r	0072	ʃ	0283
s	0073	t̪	0288
t	0074	ʋ	028A
t̪	0074 032A	ʊ	028B
tʃ	0074 0283	ʌ	028C
u	0075	ʒ	0292
u:	0075 02D0	θ	03B8

Telugu character set

For Telugu custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
-	002D	ᳵ	0C24
ç	0C01	ᳶ	0C25
o	0C02	᳷	0C26

Amazon Transcribe Developer Guide
Telugu

Character	Code	Character	Code
ః	0C03	ఁ	0C27
ఊ	0C05	ం	0C28
ఋ	0C06	ః	0C2A
ౠ	0C07	ఄ	0C2B
ఌ	0C08	అ	0C2C
఍	0C09	ఆ	0C2D
ఐ	0C0A	ఇ	0C2E
ఋ	0C0B	ఘ	0C2F
ౠ	0C0C	ఙ	0C30
ఌ	0C0E	ఠ	0C31
఍	0C0F	డ	0C32
ఐ	0C10	ఢ	0C33
ఋ	0C12	ణ	0C35
ౠ	0C13	ఠ	0C36
ఌ	0C14	డ	0C37
఍	0C15	ణ	0C38
ఐ	0C16	ఠ	0C39
ఋ	0C17	ఠ	0C3E
ౠ	0C18	ఠ	0C3F
ఌ	0C19	ఠ	0C40
఍	0C1A	ఠ	0C41
ఐ	0C1B	ఠ	0C42
ఋ	0C1C	ఠ	0C43
ౠ	0C1D	ఠ	0C44
ఌ	0C1E	ఠ	0C47
఍	0C1F	ఠ	0C48
ఐ	0C20	ఠ	0C4A
ఋ	0C21	ఠ	0C4B
ౠ	0C22	ఠ	0C4C
ఌ	0C23	ఠ	0C4D

You can use the following International Phonetic Alphabet characters in the IPA field of your input file:

Character	Code	Character	Code
ɖ	0064 032A	ð	00F0
ɖ	0064 032A 0324	ŋ	014B
dʒ	0064 0292	ɑ	0251
dʒ	0064 0292 0324	ɔ	0254
e	0065	ɔ̣	0256
e:	0065 02D0	ɔ̣̣	0256 0324
f	0066	ə	0259
h	0068	ɛ	025B
i	0069	g	0261
iẓ	0069 0290	g	0261 0324
j	006A	ɪ	026A
k	006B	ɫ	026D
k ^h	006B 02B0	ɹ	0272
l	006C	ɻ	0273
m	006D	ɻ	0279
n	006E	ɻ	0279 0329
o	006F	ɽ	027D
o:	006F 02D0	ʂ	0282
p	0070	ʃ	0283
p ^h	0070 02B0	ʈ	0288
r	0072	ʈ ^h	0288 02B0
s	0073	ʊ	028A
t	0074	ʊ	028B
ʈ	0074 032A	ʌ	028C
ʈ ^h	0074 032A 02B0	ɜ	0292
t	0074	θ	03B8
ʈ ^h	0074 0283 02B0		

Thai character set

For Thai custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
ก	0E01	ล	0E25
ข	0E02	ฬ	0E26
ฃ	0E03	ว	0E27
ค	0E04	ฅ	0E28
ฅ	0E05	๕	0E29
ฆ	0E06	ส	0E2A
ง	0E07	ห	0E2B
จ	0E08	ฬ	0E2C
ฉ	0E09	อ	0E2D
ช	0E0A	๕	0E2E
ฌ	0E0B	๗	0E2F
ฎ	0E0C	๘	0E30
ญ	0E0D	๙	0E31
ฎ	0E0E	๐	0E32
ฏ	0E0F	๑	0E34
ฐ	0E10	๒	0E35
ฑ	0E11	๓	0E36
ฒ	0E12	๔	0E37
ณ	0E13	๕	0E38
ด	0E14	๖	0E39
ต	0E15	๗	0E3A
ถ	0E16	๘	0E40
ท	0E17	๙	0E41
ธ	0E18	๐	0E42

Character	Code	Character	Code
น	0E19	ใ	0E43
บ	0E1A	ไ	0E44
ป	0E1B	ง	0E45
ผ	0E1C	จ	0E46
ฝ	0E1D	ช	0E47
พ	0E1E	ช	0E48
ฟ	0E1F	ช	0E49
ภ	0E20	ช	0E4A
ม	0E21	ช	0E4B
ย	0E22	ช	0E4C
ร	0E23	ช	0E4D
ฤ	0E24		

You can use the following International Phonetic Alphabet characters in the IPA field of your input file:

Character	Code	Character	Code
a	0061	p ^h	0070 02B0
a:	0061 02D0	r	0072
b	0062	s	0073
d	0064	t	0074
e	0065	t ^h	0074 02B0
e:	0065 02D0	t̃ɕ	0074 0361 0255
f	0066	t̃ɕ ^h	0074 0361 0255 02B0
h	0068	u	0075
i	0069	u:	0075 02D0
i:	0069 02D0	u:a	0075 02D0 0061
i:a	0069 02D0 0061	w	0077
j	006A	ŋ	014B
k	006B	ɔ	0254
kw	006B 0077	ɔ:	0254 02D0
k ^h	006B 02B0	ɛ	025B
k ^{hw}	006B 02B0 0077	ɛ:	025B 02D0

Character	Code	Character	Code
l	006C	ı	0264
m	006D	ı:	0264 02D0
n	006E	u	026F
o	006F	u:	026F 02D0
o:	006F 02D0	u:a	026F 02D0 0061
p	0070	?	0294

Turkish character set

For Turkish custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` fields:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can also use the following Unicode characters in the `Phrase` and `SoundsLike` fields:

Character	Code	Character	Code
Ç	00C7	ö	00F6
Ö	00D6	û	00FB
Ü	00DC	ü	00FC
â	00E2	ǧ	011E
ä	00E4	ǧ	011F
ç	00E7	ı	0130
è	00E8	ı	0131
é	00E9	ş	015E
ê	00EA	ş	015F
í	00ED	š	0161
î	00EE	ž	017E
ó	00F3		

You can use the following International Phonetic Alphabet characters in the `IPA` field of your input file:

Character	Code	Character	Code
a	0061	u	0075
a:	0061 02D0	u:	0075 02D0
b	0062	v	0076
c	0063	w	0077
d	0064	y	0079
e	0065	y:	0079 02D0
e:	0065 02D0	z	007A
f	0066	ø	00F8
g	0067	ø:	00F8 02D0
h	0068	ı	014B
i	0069	ı	025F
i:	0069 02D0	ı	0263
j	006A	ı	026B
k	006B	ı	026F
l	006C	ı:	026F 02D0
m	006D	r	027E
n	006E	ı	0283
o	006F	ı	0292
o:	006F 02D0	ı	0294
p	0070	ı	02A4
s	0073	ı	02A7
t	0074		

Amazon Transcribe features

Amazon Transcribe offers several services that each support a different subset of features.

In the following table, 'batch' refers to transcribing a file that is located in an Amazon S3 bucket and 'streaming' refers to transcribing media in real time.

Feature	Amazon Transcribe	Amazon Transcribe Medical (p. 256) ¹	Amazon Transcribe Call Analytics (p. 178)
<i>Configuration options</i>			
Language identification (p. 160) (single language audio)	batch, streaming	N/A	batch
Language identification (p. 160) (multi-language audio)	batch, streaming	N/A	No
Channel identification (p. 154)	batch, streaming	batch, streaming	batch
Job queueing (p. 92)	batch	No	batch
Speaker diarization (p. 172)	batch, streaming	batch, streaming	batch
Transcribing digits (p. 48) ²	batch, streaming	batch, streaming	batch
<i>Conversation analytics</i>			
Action items (p. 180) ²	No	No	batch
Call outcomes (p. 180) ²	No	No	batch
Issue detection (p. 180) ²	No	No	batch
Non-talk time (silence) (p. 180) ²	No	No	batch
Speaker interruptions (p. 180) ²	No	No	batch
Speaker sentiment (p. 181) ²	No	No	batch
Talk speed (p. 180) ²	No	No	batch
<i>Language customization</i>			
Custom language models (p. 115) ²	batch, streaming	No	batch

Feature	Amazon Transcribe	Amazon Transcribe Medical (p. 256) ¹	Amazon Transcribe Call Analytics (p. 178)
Custom vocabularies (p. 101)	batch, streaming	batch, streaming	batch
<i>Resource organization</i>			
Tagging (p. 96)	batch	batch, streaming	No
<i>Sensitive data</i>			
Identifying personal health information (p. 322)²	No	batch, streaming	No
Identifying personally identifiable information (p. 146)²	streaming	No	No
Redacting audio (p. 181)²	No	No	batch
Redacting transcripts (p. 142)²	batch, streaming	No	batch
Vocabulary filtering (p. 131)	batch, streaming	no	batch
<i>Video</i>			
Subtitles (p. 175)	batch	No	No

¹ Amazon Transcribe Medical is only available in US English.

² This feature is not available for all languages; review the [Supported languages and language-specific features \(p. 3\)](#) table for more details.

Guidelines and quotas

Amazon Transcribe has several guidelines to follow in order to achieve optimal results. There are also quotas that may impact your transcriptions; some of these can be increased. Refer to the following sections for details.

Supported AWS Regions

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

Guidelines

For best results:

- Use a lossless format, such as FLAC or WAV with PCM 16-bit encoding.
- Use a sample rate of 8,000 Hz for low-fidelity audio and 16,000-48,000 Hz for high-fidelity audio.

If you don't need to process all of your transcription jobs concurrently, use [Job queueing \(p. 92\)](#). This enables Amazon Transcribe to keep track of your transcription jobs and process them when slots are available.

Note

Amazon Transcribe may temporarily store your content to continuously improve the quality of its analysis models. See the [Amazon Transcribe FAQ](#) to learn more. To request the deletion of content that may have been stored by Amazon Transcribe, open a case with [AWS Support](#).

Quotas

The following quotas **cannot** be changed:

Description	Quota
Audio file length	4:00:00 (four) hours (14,400 seconds)
Audio stream duration	4:00:00 (four) hours (14,400 seconds)
Audio file size	2 GB
Audio file size (call analytics)	500 MB
Size of a custom vocabulary	50 KB
Length of a custom vocabulary phrase	256 characters
Size of a custom vocabulary filter	50 KB

Description	Quota
Number of custom vocabulary filters	100
Number of channels for channel identification	2
Number of days job records are retained	90
Minimum audio file duration	500 milliseconds (ms)

You can request a quota increase for the following resources:

Resource	Default
Number of concurrent batch transcription jobs	250
Number of concurrent batch transcription jobs (Call Analytics)	100
Job queue bandwidth ratio	0.9
Number of concurrent HTTP/2 requests	25
Number of concurrent WebSocket requests	25
Total number of vocabularies per AWS account	100
Number of pending vocabularies	10
Number of concurrently training custom language models	3
Total number of custom language models per AWS account	10
Number of categories per AWS account (Call Analytics)	200
Number of rules per category (Call Analytics)	20

The following operations limits can also be increased upon request:

Operation	Transactions per second
StartTranscriptionJob	25

Operation	Transactions per second
StartStreamTranscription	25
GetTranscriptionJob	30
DeleteTranscriptionJob	5
ListTranscriptionJobs	5
CreateVocabulary	10
UpdateVocabulary	10
DeleteVocabulary	5
GetVocabulary	20
ListVocabularies	5
StartCallAnalyticsJob	10
GetCallAnalyticsJob	20
ListCallAnalyticsJobs	5
DeleteCallAnalyticsJob	5
CreateCallAnalyticsCategory	10
UpdateCallAnalyticsCategory	10
DeleteCallAnalyticsCategory	5
GetCallAnalyticsCategory	20
ListCallAnalyticsCategories	5

Note

For information about requesting a quota increase, see [AWS service quotas](#) in the *AWS General Reference*.

How Amazon Transcribe works

Amazon Transcribe converts speech to text. A basic transcription request produces a transcript that contains data about the transcribed content, including confidence scores and timestamps for each word or punctuation mark. For a complete list of features that you can apply to your transcription, refer to the [feature summary](#) (p. 40).

Transcription methods can be separated into two main categories:

- **Batch transcription jobs:** Transcribe media files that have been uploaded into an Amazon S3 bucket.
- **Streaming transcriptions:** Transcribe media streams in real time.

You can create batch transcriptions using the [AWS CLI](#) (p. 63), [AWS Management Console](#) (p. 55), and various [AWS SDKs](#) (p. 66).

You can create streaming transcriptions using the [AWS Management Console](#) (p. 55), [HTTP/2](#) (p. 82), [WebSockets](#) (p. 86), and various [AWS SDKs](#) (p. 66).

Topics

- [Data input and output](#) (p. 45)
- [Transcribing numbers and punctuation](#) (p. 48)

API operations to get you started

Batch: [StartTranscriptionJob](#)

Streaming: [StartStreamTranscription](#)

Data input and output

Amazon Transcribe takes audio data, either a media file in an Amazon S3 bucket or a media stream, and converts it to text data.

If you're working with media files stored in an Amazon S3 bucket, you're performing batch transcription jobs; if you're working with media streams, you're performing streaming transcriptions. These two processes have different rules and requirements.

Media formats

Supported media types differ between batch transcriptions and streaming transcriptions, though lossless formats are recommended for both. See the following table for details:

	Batch	Streaming
Supported formats	<ul style="list-style-type: none">• AMR• FLAC• MP3• MP4• Ogg• WebM• WAV	<ul style="list-style-type: none">• FLAC• OPUS-encoded audio in an Ogg container• PCM signed 16-bit little-endian audio (note that this does not include WAV)

	Batch	Streaming
Recommended formats	<ul style="list-style-type: none">• FLAC• WAV with PCM 16-bit encoding	<ul style="list-style-type: none">• FLAC• PCM encoding

Sample rates

With batch transcription jobs, you can choose to provide a sample rate, though this parameter is optional. If you do include it in your request, make sure the value you provide matches the actual sample rate in your audio; if you provide a sample rate that doesn't match your audio, your job may fail.

With streaming transcriptions, you must include a sample rate in your request. As with batch transcription jobs, make sure the value you provide matches the actual sample rate in your audio.

Sample rates for low fidelity audio, such as telephone recordings, typically use 8,000 Hz. For high fidelity audio, Amazon Transcribe supports values between 16,000 Hz and 48,000 Hz.

Output

Transcription output is in JSON format. The first part of your transcript contains the transcript itself, in paragraph form, followed by additional data for every word and punctuation mark. The data provided depends on the features you include in your request. At a minimum, your transcript contains the start time, end time, and confidence score for every word. The [following section \(p. 46\)](#) shows example output from a basic transcription request that didn't include any additional options or features.

All **batch transcripts** are stored in Amazon S3 buckets. You can either choose to save your transcript in your own Amazon S3 bucket, or have Amazon Transcribe use a secure default bucket. To learn more about creating and using Amazon S3 buckets, see [Working with buckets](#).

If you want your transcript stored in an Amazon S3 bucket that you own, specify the bucket's URI in your transcription request. Make sure you give Amazon Transcribe write permissions for this bucket prior to starting your batch transcription job. If you specify your own bucket, your transcript remains in that bucket until you remove it.

If you don't specify an Amazon S3 bucket, Amazon Transcribe uses a secure service-managed bucket and provides you with a temporary URI you can use to download your transcript. Note that temporary URIs are valid for 15 minutes. If you get an `AccessDenied` error when using the provided URI, make a `GetTranscriptionJob` request to get a new temporary URI for your transcript.

If you opt for a default bucket, your transcript is deleted when your job expires (90 days). If you want to keep your transcript past this expiration date, you must download it.

Streaming transcripts are returned via the method you're using for your stream (WebSocket or HTTP/2).

Tip

If you want to convert your JSON output into a turn-by-turn transcript in Word format, see this [GitHub example \(for Python3\)](#).

Example output

Transcripts provide you with a complete transcription in paragraph form, followed by a word-for-word breakdown, which provides data for every word and punctuation mark. This includes start time, end time, a confidence score, and a type (either `pronunciation` or `punctuation`). The below example is

a transcript from a simple transcription request that didn't include any [additional features \(p. 40\)](#). With each additional feature that you apply to your transcription request, you get additional data in your transcript output file.

```
{
  "jobName": "my-first-transcription-job",
  "accountId": "111122223333",
  "results": {
    "transcripts": [
      {
        "transcript": "Welcome to Amazon Transcribe."
      }
    ],
    "items": [
      {
        "start_time": "0.64",
        "end_time": "1.09",
        "alternatives": [
          {
            "confidence": "1.0",
            "content": "Welcome"
          }
        ],
        "type": "pronunciation"
      },
      {
        "start_time": "1.09",
        "end_time": "1.21",
        "alternatives": [
          {
            "confidence": "1.0",
            "content": "to"
          }
        ],
        "type": "pronunciation"
      },
      {
        "start_time": "1.21",
        "end_time": "1.74",
        "alternatives": [
          {
            "confidence": "1.0",
            "content": "Amazon"
          }
        ],
        "type": "pronunciation"
      },
      {
        "start_time": "1.74",
        "end_time": "2.56",
        "alternatives": [
          {
            "confidence": "1.0",
            "content": "Transcribe"
          }
        ],
        "type": "pronunciation"
      },
      {
        "alternatives": [
          {
            "confidence": "0.0",
            "content": "."
          }
        ]
      }
    ]
  }
}
```

```

        "type": "punctuation"
      }
    ],
    "status": "COMPLETED"
  }

```

Transcribing numbers and punctuation

Amazon Transcribe automatically adds punctuation to all supported languages, and capitalizes words appropriately for languages that use case distinction in their writing systems.

For most languages, numbers are transcribed into their word forms. However, if your media is in English or German, Amazon Transcribe treats numbers differently depending on the context in which they're used.

For example, if a speaker says *"Meet me at eight-thirty AM on June first at one-hundred Main Street with three-dollars-and-fifty-cents and one-point-five chocolate bars,"* this is transcribed as:

- English and German dialects: Meet me at 8:30 a.m. on June 1st at 100 Main Street with \$3.50 and 1.5 chocolate bars
- All other languages: Meet me at eight thirty a m on June first at one hundred Main Street with three dollars and fifty cents and one point five chocolate bars

To see all the rules associated with spoken numbers in English and German, refer to the following table.

Rules	English dialects (input audio → output text)	German dialects (input audio → output text)
Convert cardinal numbers greater than ten to numbers.	<ul style="list-style-type: none"> • "Fifty five" → 55 • "a hundred" → 100 • "One thousand and thirty one" → 1031 • "One hundred twenty-three million four hundred fifty six thousand seven hundred eight nine" → 123,456,789 	<ul style="list-style-type: none"> • "fünfundfünfzig" → 55 • "vier tausend sechs hundert einundachtzig" → 4681 • "eine Sache" → "eine Sache"
Convert cardinal numbers followed by "million" or "billion" to numerals followed by a word when "million" or "billion" is not followed by a number.	<ul style="list-style-type: none"> • "one hundred million" → 100 million • "one billion" → 1 billion • "two point three million" → 2.3 million 	<ul style="list-style-type: none"> • "zehn Millionen Menschen" → 10 Millionen Menschen • "zehn Millionen fünf hundert tausend" → 10.500.000
Convert ordinal numbers greater than ten to numbers.	<ul style="list-style-type: none"> • "Forty third" → 43rd • "twenty sixth avenue" → 26th avenue 	<ul style="list-style-type: none"> • "dreiundzwanzigste" → 23 • "vierzigster" → 40 • "ich war Erster" → "ich war Erster"
Convert fractions to their numeric format.	<ul style="list-style-type: none"> • "a quarter" → 1/4 • "three sixteenths" → 3/16 • "a half" → 1/2 	Fractions are not converted into a numeric format. <ul style="list-style-type: none"> • "ein Drittel" → "ein Drittel"

Rules	English dialects (input audio → output text)	German dialects (input audio → output text)
	<ul style="list-style-type: none"> "a hundredth" → 1/100 	
Convert numbers less than ten to digits if there are more than one in a row.	<ul style="list-style-type: none"> "three four five" → 345 "My phone number is four two five five five five one two one two" → My phone number is 4255551212 	<ul style="list-style-type: none"> "eins zwei drei" → 123 "plus vier neun zwei vier eins" → +49241
The words "dot" or "point" are displayed as a decimal.	<ul style="list-style-type: none"> "three hundred and three dot five" → 303.5 "three point twenty three" → 3.23 "zero point four" → 0.4 "point three" → 0.3 	Decimals are indicated by ",". <ul style="list-style-type: none"> "zweiundzwanzig komma drei" → 22,3
Convert the word "percent" after a number to a percent symbol (%).	<ul style="list-style-type: none"> "twenty three percent" → 23% "twenty three point four five percent" → 23.45% 	<ul style="list-style-type: none"> "fünf Prozent Hürde" → 5% Hürde "dreiundzwanzig komma vier Prozent" → 23,4%

Rules	English dialects (input audio → output text)	German dialects (input audio → output text)
Convert monetary words to symbols.	<p>Convert the words "dollar," "U S dollar," "Australian dollar," "AUD," or "USD" after a number to a dollar sign (\$) before the number.</p> <ul style="list-style-type: none"> "one dollar and fifteen cents" → \$1.15 "twenty three USD" → \$23 "twenty three Australian dollars" → \$23 <p>Convert the words "pounds," "British pounds," or "GDB" after a number to pound sign (£) before the number.</p> <ul style="list-style-type: none"> "twenty three pounds" → £23 "I have two thousand pounds" → I have £2,000 "five pounds thirty three pence" → £5.33 <p>Convert the words "rupees," "Indian rupees," or "INR" after a number to rupee sign (₹) before the number.</p> <ul style="list-style-type: none"> "twenty three rupees" → ₹23 "fifty rupees thirty paise" → ₹50.30 	<p>Convert the words "Euro" to a euro sign.</p> <ul style="list-style-type: none"> "ein euro" → 1 € "ein Euro vierzig" → 1,40 € "ein Euro vierzig Cent" → 1,40 €
Convert times to numbers.	<ul style="list-style-type: none"> "seven a m eastern standard time" → 7 a.m. eastern standard time "twelve thirty p m" → 12:30 p.m. 	<ul style="list-style-type: none"> "vierzehn Uhr fünfzehn" → 14:15 Uhr
Convert dates to numbers.	<ul style="list-style-type: none"> "May fifth twenty twelve" → May 5th 2012 "May five twenty twelve" → May 5 2012 "five May twenty twelve" → 5 May 2012 	<ul style="list-style-type: none"> "dritter Dezember neunzehn hundert sechundfünfzig" → 3. Dezember 1956
Separate spans of numbers by the word "to".	<ul style="list-style-type: none"> "twenty three to thirty seven" → 23 to 37 	Not applicable

Rules	English dialects (input audio → output text)	German dialects (input audio → output text)
Years are represented as four digits; this is only valid for years in the 20th, 21st, and 22nd centuries.	<ul style="list-style-type: none"> • "nineteen sixty two" → 1962 • "the year is twenty twelve" → the year is 2012 • "twenty nineteen" → 2019 • "twenty one thirty" → 2130 	Not applicable
Display slashes and dashes.	<ul style="list-style-type: none"> • "fifty-five dash thirteen" → 55-13 <p>Slashes are not displayed.</p> <ul style="list-style-type: none"> • "fifty-five slash thirteen" → 55 slash 13 	<ul style="list-style-type: none"> • "fünfundfünfzig Schrägstrich dreizehn" → 55/13 • "fünfundfünfzig Strich dreizehn" → 55-13
Display numbered paragraphs.	<p>Numbered paragraphs are not displayed using the paragraph symbol (§).</p> <ul style="list-style-type: none"> • "paragraph seventeen" → paragraph 17 	<ul style="list-style-type: none"> • "Paragraf siebzehn" → § 17

Getting started with Amazon Transcribe

Before you can create transcriptions, you have a few prerequisites:

- [Sign up for an AWS account \(p. 52\)](#)
- [Install the AWS CLI and SDKs \(p. 52\)](#) (if you're using the AWS Management Console for your transcriptions, you can skip this step)
- [Create an IAM user \(p. 53\)](#) with administrator permissions
- [Set up an Amazon S3 bucket \(p. 53\)](#)
- [Create an IAM policy \(p. 53\)](#)

Once you complete these prerequisites, you're ready to transcribe. Select your preferred transcription method from the following list to get started.

- [AWS CLI \(p. 63\)](#)
- [AWS Management Console \(p. 55\)](#)
- [AWS SDK \(p. 66\)](#)
- [HTTP \(p. 74\)](#)
- [WebSockets \(p. 74\)](#)

Because streaming using HTTP/2 and WebSockets is more complicated than the other transcription methods, we advise reviewing the [Setting up a streaming transcription \(p. 80\)](#) section before getting started with these methods.

Signing up for an AWS account

You can sign up for either a [free tier](#) account or a [paid account](#). Both options give you access to all AWS services. The free tier has a trial period during which you can explore AWS services and estimate your usage. Once your trial period expires, you can migrate to a paid account. Fees are accrued on a pay-as-you-use basis; see [Amazon Transcribe Pricing](#) for details.

Tip

When setting up your account, make note of your AWS account ID because you need it to create an IAM user or group.

Installing the AWS CLI and SDKs

To use the Amazon Transcribe API, you must first install the AWS CLI. The current AWS CLI is version 2. You can find installation instructions for [Linux](#), [Mac](#), [Windows](#), and [Docker](#) in the [AWS Command Line Interface User Guide](#).

Once you have the AWS CLI installed, you need to [configure](#) it for your security credentials and AWS Region.

If you want to use Amazon Transcribe with an SDK, select your preferred language for installation instructions:

- [.NET](#)
- [C++](#)
- [Go](#)
- [Java V2](#)
- [JavaScript](#)
- [PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#) (batch transcriptions only)
- [Python](#) (streaming transcriptions only)
- [Ruby V3](#)

Creating an IAM user

IAM, or Identity and Access Management, is a means of controlling which users can access which resources. All AWS services use IAM credentials and it is best practice to use an IAM admin user to access your resources. To learn more about IAM, see [What is IAM?](#)

You can create an IAM admin user with the AWS Management Console or AWS CLI. For instructions, see [Creating your first IAM admin user and user group](#). For information on signing in to your account using IAM user credentials, refer to [How IAM users sign in to your AWS account](#).

Using your AWS account root user credentials to access resources, such as a transcription request, is **not** recommended. Creating non-admin identities with specific (restricted) permissions is best practice. For instructions on creating non-admin identities, see [Creating your first IAM delegated user and user group](#).

Creating an Amazon S3 bucket

Amazon S3 is a secure object storage service. Amazon S3 stores your files (called *objects*) in containers (called *buckets*).

To run a batch transcription, you must first upload your media files into an Amazon S3 bucket. If you don't specify an Amazon S3 bucket for your transcription output, Amazon Transcribe puts your transcript in a temporary AWS-managed Amazon S3 bucket. Transcription output in AWS-managed buckets is automatically deleted after 90 days.

Learn how to [Create your first S3 bucket](#) and [Upload an object to your bucket](#).

Creating an IAM policy

To manage access in AWS, you must create policies and attach them to IAM identities (users, groups, or roles) or AWS resources. A policy defines the permissions of the entity it is attached to. For example, a user can only access a media file located in your Amazon S3 bucket if you've attached a policy to that user which grants them access. If you want to further restrict that user, you can instead limit their access to a specific file within an Amazon S3 bucket.

To learn more about using AWS policies see:

- [Policies and permissions in IAM](#)
- [Creating IAM policies](#)
- [How Amazon Transcribe works with IAM \(p. 229\)](#)

For example policies you can use with Amazon Transcribe, see [Amazon Transcribe identity-based policy examples \(p. 235\)](#). If you want to generate custom policies, consider using the [AWS Policy Generator](#).

You can add a policy using the AWS Management Console, AWS CLI, or AWS SDK. For instructions, see [Adding and removing IAM identity permissions](#).

Policies have the format:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "my-policy-name",
      "Effect": "Allow",
      "Action": [
        "service:action"
      ],
      "Resource": [
        "amazon-resource-name"
      ]
    }
  ]
}
```

Amazon Resource Names (ARNs) uniquely identify all AWS resources, such as an Amazon S3 bucket. You can use ARNs in your policy to grant permissions for specific actions to use specific resources. For example, if you want to grant read access to an Amazon S3 bucket and its sub-folders, you can add the following code to your trust policy's Statement section:

```
{
  "Effect": "Allow",
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
  ]
}
```

Here's an example policy that grants Amazon Transcribe read (GetObject, ListBucket) and write (PutObject) permissions to an Amazon S3 bucket, `DOC-EXAMPLE-BUCKET`, and its sub-folders:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "my-policy-name",
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "transcribe.amazonaws.com"
        ]
      }
    }
  ]
}
```

```
    ],  
    },  
    "Action": "sts:AssumeRole"  
  },  
  {  
    "Effect": "Allow",  
    "Action": [  
      "s3:GetObject",  
      "s3:ListBucket"  
    ],  
    "Resource": [  
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET",  
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"  
    ]  
  },  
  {  
    "Effect": "Allow",  
    "Action": [  
      "s3:PutObject",  
    ],  
    "Resource": [  
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET",  
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"  
    ]  
  }  
]  
}
```

Transcribing with the AWS Management Console

Prior to starting a batch transcription, you must first upload your media file to an Amazon S3 bucket. For streaming transcriptions using the AWS Management Console, you must use your computer microphone and transcribe real-time speech.


To view supported media formats and other media requirements and constraints, see [Data input and output \(p. 45\)](#).

Expand the following sections for short walkthroughs of each transcription method.

Batch transcriptions

First make sure that you've uploaded the media file you want to transcribe into an Amazon S3 bucket. If you're unsure how to do this, refer to the Amazon S3 User Guide: [Upload an object to your bucket](#).

1. From the [AWS Management Console](#), select **Transcription jobs** in the left navigation pane. This takes you to a list of your transcription jobs.



2. Select **Create job**.
2. Complete the fields on the **Specify job details** page.

Specify job details [Info](#)

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Model type [Info](#)
Choose the type of model to use for the transcription job.

General model
To use a model that is not specialized for a particular use case, choose this option. Configuration options vary between languages.

Custom language model
To use a model that you trained for your specific use case, choose this option. This model has fewer configuration options than the general model.

Language settings
You can transcribe your audio file in a language that you specify or have Amazon Transcribe Identify and transcribe it in the predominant language.

Specific language [Info](#)
If you know the language spoken in your source audio, choose this option to get the most accurate results. The options available for additional processing vary between languages.

Automatic language identification [Info](#)
If you don't know the language spoken in your audio files, choose this option. You have access to fewer options for additional processing than if you choose **Specific language**.

Language
Choose the language of the input audio.

▶ **Additional settings**

The input location *must* be an object within an Amazon S3 bucket. For output location, you can choose a secure Amazon S3 service-managed bucket or you can specify your own Amazon S3 bucket.

If you choose a service-managed bucket, you can view a transcript preview in the AWS Management Console and you can download your transcript from the job details page (see below).

If you choose your own Amazon S3 bucket, you cannot see a preview in the AWS Management Console and must go to the Amazon S3 bucket to download your transcript.

The screenshot shows the 'Configure job' page in the AWS Management Console. It is divided into three main sections:

- Input data**: Includes a section for 'Input file location on S3' with a text input field containing 's3://bucket/prefix/file.mp3' and a 'Browse S3' button. Below this is a note: 'Valid file formats: MP3, MP4, WAV, FLAC, AMR, OGG, and WebM.'
- Output data**: Includes a section for 'Output data location type' with two radio button options: 'Service-managed S3 bucket' (selected) and 'Customer specified S3 bucket'. Below this is a section for 'Subtitle file format' with two checkboxes: 'SRT (SubRip)' and 'VTT (WebVTT)'. Both are currently unchecked.
- Tags - optional**: Includes a description of tags, a note that there are 'No tags associated with the resource', and an 'Add new tag' button. A note below the button states: 'You can add up to 50 more tags.'

At the bottom right of the form, there are two buttons: 'Cancel' and 'Next'.

Select **Next**.

3. Select any desired options on the **Configure job** page. If you want to use [Custom vocabularies \(p. 101\)](#) or [Custom language models \(p. 115\)](#) with your transcription, you must create these before starting your transcription job.

Configure job - *optional* [Info](#)

Audio settings

Audio identification [Info](#)
Choose to split multi-channel audio into separate channels for transcription, or identify speakers in the input audio.

Alternative results [Info](#)
Enable to view more transcription results

Content removal

Content removal conceals information in the resulting transcript from your source audio file. Amazon Transcribe changes items in the transcript and does not modify the source audio.

Automatic content redaction [Info](#)
Automatic content redaction removes personally identifiable information (PII) in your transcripts. Redactions in transcripts show up as [PII].

Vocabulary filtering [Info](#)
Vocabulary filtering can remove, mask or tag specified words in the final transcript.

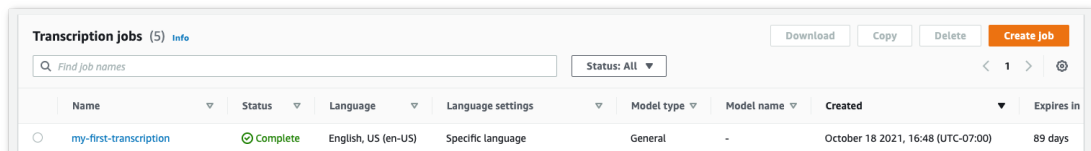
Customization

Custom vocabulary [Info](#)
A custom vocabulary improves the accuracy of recognizing words and phrases specific to your use case.

[Cancel](#) [Previous](#) [Create job](#)

Select **Create job**.

- You're now on the **Transcription jobs** page. Here you can see the status of the transcription job. Once complete, click on the name of your transcription.



Transcription jobs (5) [Info](#) [Download](#) [Copy](#) [Delete](#) [Create job](#)

[Status: All](#) [<](#) [1](#) [>](#) [⊗](#)

Name	Status	Language	Language settings	Model type	Model name	Created	Expires in
my-first-transcription	Complete	English, US (en-US)	Specific language	General	-	October 18 2021, 16:48 (UTC-07:00)	89 days

- You're now viewing the **Job details** page for your transcription. Here you can view all of the options you specified when setting up your transcription job.

To view your transcript, select the linked filepath in the right column under **Output data location**. This takes you to the Amazon S3 output folder you specified. Click on the name you of your output file, which now has a .json extension.

The screenshot displays the 'Job details' section for a transcription job named 'my-first-transcription-job'. The interface includes 'Delete' and 'Copy' buttons in the top right corner. The details are organized into four columns:

Job details			
Name	Model	Audio identification	Input data location
my-first-transcription-job	None	Disabled	s3://DOC-EXAMPLE-BUCKET/my-media-file.flac
Status	Created	Alternative results	Output data location
Complete	2/4/2022, 12:11:31 PM	Disabled	Service-managed S3 bucket
Language	Started	Custom vocabulary	
English, US (en-US)	2/4/2022, 12:11:31 PM	None	
Language settings	Ended	PII redaction	
Specific language	2/4/2022, 12:12:50 PM	Disabled	
Expiration Info	Input file format	Vocabulary filter	
The transcription is available for 87 more days.	flac	-	
	Audio sampling rate		
	44100 Hz		

6. How you download your transcript depends on whether you chose a service-managed Amazon S3 bucket or your own Amazon S3 bucket.
 - a. If you chose a service-managed bucket, you can see a **Transcription preview** pane on your transcription job's information page, along with a **Download** button.

The screenshot displays the AWS Management Console interface for a transcription job named "my-first-transcription-job". At the top right, there are "Delete" and "Copy" buttons. The main content is divided into two sections: "Job details" and "Transcription preview".

Job details

Name my-first-transcription-job	Model None	Audio identification Disabled	Input data location s3://DOC-EXAMPLE-BUCKET/my-media-file.flac
Status Complete	Created 2/4/2022, 12:11:31 PM	Alternative results Disabled	Output data location Service-managed S3 bucket
Language English, US (en-US)	Started 2/4/2022, 12:11:31 PM	Custom vocabulary None	
Language settings Specific language	Ended 2/4/2022, 12:12:50 PM	PII redaction Disabled	
Expiration Info The transcription is available for 87 more days.	Input file format flac	Vocabulary filter -	
	Audio sampling rate 44100 Hz		

Transcription preview

You can see the first 5,000 characters of the transcription text below. To download the full text, choose [Download full transcript](#).

[Download](#)

Text | **Audio identification** | **Subtitles**

This is a preview of the content of your transcript. If your transcript is long, you may have to scroll to see the complete preview.

Select **Download** and choose **Download transcript**.

- b. If you chose your own Amazon S3 bucket, you don't see any text in the **Transcription preview** pane on your transcription job's information page. Instead, you see a blue information box with a link to the Amazon S3 bucket you chose.

The screenshot displays the 'Job details' section for a transcription job. The job name is 'my-first-transcription-job'. The status is 'Complete'. The language is 'English, US (en-US)'. The job was created on 2/7/2022 at 11:42:17 AM and ended at 11:43:37 AM. The input data location is 's3://DOC-EXAMPLE-BUCKET/my-media-file.flac'. The output data location is 'https://s3.us-west-2.amazonaws.com/DOC-EXAMPLE-BUCKET'. The audio identification is disabled, alternative results are disabled, and PII redaction is disabled. The input file format is 'flac' and the audio sampling rate is '44100 Hz'. There is also a 'Transcription preview' section with a 'Download' button and a note about S3 bucket output.

Job details			
Name	Model	Audio identification	Input data location
my-first-transcription-job	None	Disabled	s3://DOC-EXAMPLE-BUCKET/my-media-file.flac
Status	Created	Alternative results	Output data location
Complete	2/7/2022, 11:42:17 AM	Disabled	https://s3.us-west-2.amazonaws.com/DOC-EXAMPLE-BUCKET
Language	Started	Custom vocabulary	
English, US (en-US)	2/7/2022, 11:42:17 AM	None	
Language settings	Ended	PII redaction	
Specific language	2/7/2022, 11:43:37 AM	Disabled	
Expiration	Input file format	Vocabulary filter	
The transcription is available for 89 more days.	flac	-	
	Audio sampling rate		
	44100 Hz		

To access your transcript, go to the specified Amazon S3 bucket using either the link under **Output data location** in the **Job details** pane or the **S3 Bucket** link within the blue information box in the **Transcription preview** pane.

Streaming transcriptions

1. From the **AWS Management Console**, select **Real-time transcription** in the left navigation pane. This takes you to the main streaming page where you can select options prior to starting your stream.

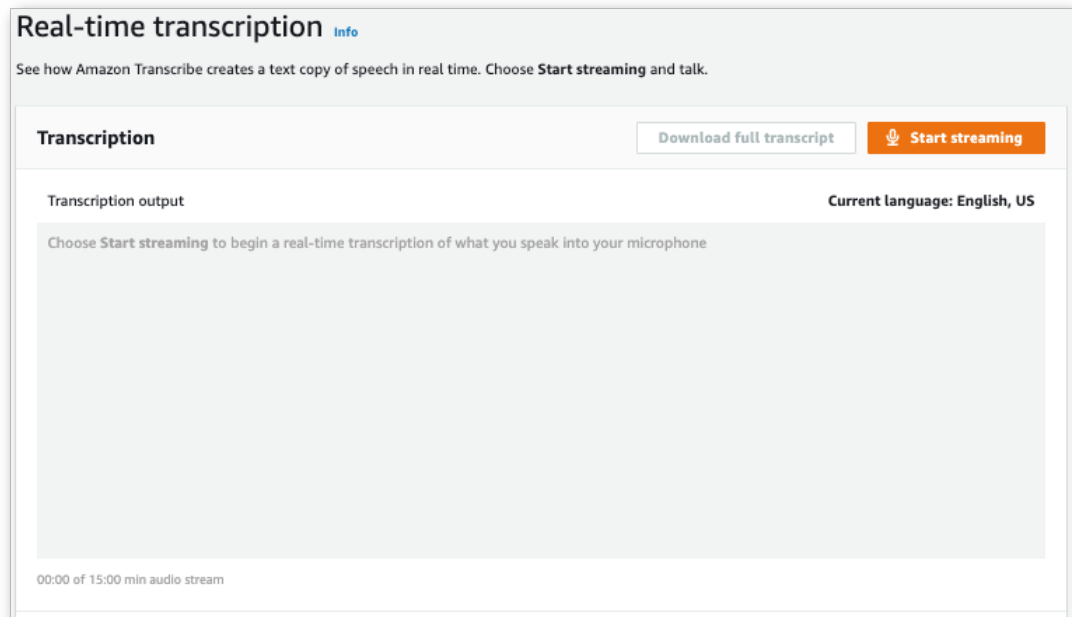
The screenshot shows the 'Real-time transcription' page in the AWS Management Console. The left navigation pane is open, showing 'Real-time transcription' selected. The main content area has a 'Transcription' section with a 'Download full transcript' button and a 'Start streaming' button. Below this is a 'Transcription output' section with a text area for transcription results and a 'Current language: English, US' indicator.

2. Below the **Transcription output** box, you have the option to select various language and audio settings.

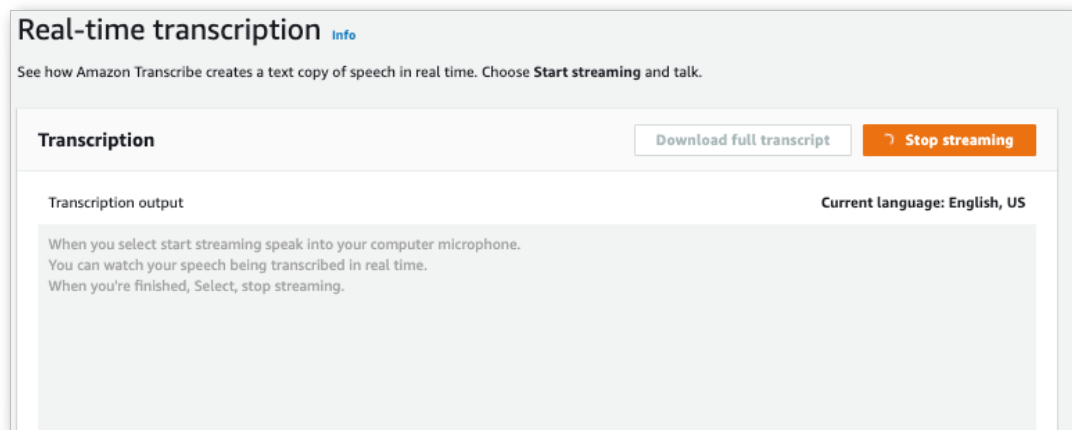
The screenshot shows the settings panel for Amazon Transcribe. It is organized into several sections, each with a dropdown arrow on the left:

- Language settings**
 - Language settings**: You can select a specific language for your transcription or have Amazon Transcribe identify the predominant language in your media and perform the transcription in that language.
 - Specific language**: If you know the language spoken in your source audio, choose this option to get the most accurate results.
 - Automatic language identification** [Info](#): If you don't know the language spoken in your audio files, choose this option.
 - Language**: A dropdown menu currently set to "English, US (en-US)".
- Audio settings**
 - Speaker identification** [Info](#): Identify the different speakers in the stream. Speaker identification might vary in availability between languages.
- Content removal settings**
 - Vocabulary filtering** [Info](#): Vocabulary filtering removes, masks, or tags words that you specify in your vocabulary filter. Choose a vocabulary filter to see an example.
 - PII Identification & redaction** [Info](#): Identify or redact one or more types of personally identifiable information (PII) in your transcript.
- Customizations**
 - Custom vocabulary** [Info](#): A custom vocabulary improves the accuracy of recognizing words and phrases specific to your use case.
 - Partial results stabilization** [Info](#): Configure Amazon Transcribe to present results that don't change as it processes the transcription output from your stream.
 - Custom language model** [Info](#): Select the model you want to use with this streaming session.

3. After you've selected the appropriate settings, scroll to the top of the page and choose **Start streaming**, then begin speaking into your computer microphone. You can see your speech transcribed in real time.



4. When you're finished, select **Stop streaming**.



You can now download your transcript by selecting **Download full transcript**.

Transcribing with the AWS CLI

When using the AWS CLI to start a transcription, you can either run all commands at the CLI level or you can run the Amazon Transcribe command you wish to use followed by the AWS Region and the location of a JSON file that contains a request body. Examples throughout this guide show both methods; however, this section focuses on the former method.

The AWS CLI does not support streaming transcriptions.

Before you continue, make sure you've:

- Uploaded your media file into an Amazon S3 bucket. If you're unsure how to create an Amazon S3 bucket or upload your file, refer to [Create your first Amazon S3 bucket](#) and [Upload an object to your bucket](#).
- Installed the [AWS CLI \(p. 52\)](#).

You can find all AWS CLI commands for Amazon Transcribe in the [AWS CLI Command Reference](#).

Starting a new transcription job

To start a new transcription, use the `start-transcription-job` command.

1. In a terminal window, type the following:

```
aws transcribe start-transcription-job \
```

A `>` appears on the next line and you can now continue adding required parameters, as described in the next step.

You can also omit the `\` and append all parameters, separating each with a space.

2. With the `start-transcription-job` command, you must include `region`, `transcription-job-name`, `media`, and either `language-code` or `identify-language`.

If you want to specify an output location, include `output-bucket-name` in your request; if you want to specify a sub-folder of the specified output bucket, also include `output-key`.

```
aws transcribe start-transcription-job \  
  --region us-west-2 \  
  --transcription-job-name my-first-transcription-job \  
  --media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
  --language-code en-US
```

If appending all parameters, this request looks like:

```
aws transcribe start-transcription-job --region us-west-2 --transcription-job-name my-first-transcription-job --media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac --language-code en-US
```

If you choose not to specify an output bucket using `output-bucket-name`, Amazon Transcribe places your transcription output in a service-managed bucket. Transcripts stored in a service-managed bucket expire after 90 days.

Amazon Transcribe responds with:

```
{  
  "TranscriptionJob": {  
    "TranscriptionJobName": "my-first-transcription-job",  
    "TranscriptionJobStatus": "IN_PROGRESS",  
    "LanguageCode": "en-US",  
    "Media": {  
      "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
    },  
    "StartTime": "2022-03-07T15:03:44.246000-08:00",  
    "CreationTime": "2022-03-07T15:03:44.229000-08:00"  
  }  
}
```

To see if your transcription job is successful, as indicated by `TranscriptionJobStatus` changing from `IN_PROGRESS` to `COMPLETED`, use the `get-transcription-job` or `list-transcription-job` command, as shown in the following section.

Getting the status of a transcription job

To get information about your transcription job, use the `get-transcription-job` command.

The only required parameters for this command are the AWS Region where the job is located and the name of the job.

```
aws transcribe get-transcription-job \  
  --region us-west-2 \  
  --transcription-job-name my-first-transcription-job
```

Amazon Transcribe responds with:

```
{  
  "TranscriptionJob": {  
    "TranscriptionJobName": "my-first-transcription-job",  
    "TranscriptionJobStatus": "COMPLETED",  
    "LanguageCode": "en-US",  
    "MediaSampleRateHertz": 48000,  
    "MediaFormat": "flac",  
    "Media": {  
      "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
    },  
    "Transcript": {  
      "TranscriptFileUri": "https://s3.the-URI-where-your-job-is-located.json"  
    },  
    "StartTime": "2022-03-07T15:03:44.246000-08:00",  
    "CreationTime": "2022-03-07T15:03:44.229000-08:00",  
    "CompletionTime": "2022-03-07T15:04:01.158000-08:00",  
    "Settings": {  
      "ChannelIdentification": false,  
      "ShowAlternatives": false  
    }  
  }  
}
```

If you've selected your own Amazon S3 bucket for your transcription output, this bucket is listed with `TranscriptFileUri`. If you've selected a service-managed bucket, a temporary URI is provided; use this URI to download your transcript.

Note

Temporary URIs for service-managed Amazon S3 buckets are only valid for 15 minutes. If you get an `AccessDenied` error when using the URI, run the `get-transcription-job` request again to get a new temporary URI.

Listing your transcription jobs

To list all your transcription jobs in a given AWS Region, use the `list-transcription-jobs` command.

The only required parameter for this command is the AWS Region in which your transcription jobs are located.

```
aws transcribe list-transcription-jobs \  
  --region us-west-2
```

Amazon Transcribe responds with:

```
{
  "NextToken": "A-very-long-string",
  "TranscriptionJobSummaries": [
    {
      "TranscriptionJobName": "my-first-transcription-job",
      "CreationTime": "2022-03-07T15:03:44.229000-08:00",
      "StartTime": "2022-03-07T15:03:44.246000-08:00",
      "CompletionTime": "2022-03-07T15:04:01.158000-08:00",
      "LanguageCode": "en-US",
      "TranscriptionJobStatus": "COMPLETED",
      "OutputLocationType": "SERVICE_BUCKET"
    }
  ]
}
```

Deleting your transcription job

To delete your transcription job, use the `delete-transcription-job` command.

The only required parameters for this command are the AWS Region where the job is located and the name of the job.

```
aws transcribe delete-transcription-job \
  --region us-west-2 \
  --transcription-job-name my-first-transcription-job
```

To confirm your delete request is successful, you can run the `list-transcription-jobs` command; your job should no longer appear in the list.

Transcribing with the AWS SDKs

You can use SDKs for both batch and streaming transcriptions.

To get started with batch transcriptions, refer to [Batch transcriptions \(p. 66\)](#). Note that prior to starting a batch transcription, you must first upload your media file into an Amazon S3 bucket.

To get started with streaming transcriptions, refer to [Streaming transcriptions \(p. 69\)](#). For streaming transcriptions, you can transcribe a streamed media file or a live media stream. Note that streaming transcriptions are not supported with the following AWS SDKs: Java V1, .NET, Python, or PowerShell. If you wish to use Python for streaming, there's a custom [Python SDK for Amazon Transcribe](#) available.

To view supported media formats and other media requirements and constraints, see [Data input and output \(p. 45\)](#).

Tip

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

You can also find code samples in the [AWS Samples repository](#) on GitHub.

Batch transcriptions

You can create batch transcriptions using the URI of a media file located in an Amazon S3 bucket. If you're unsure how to create an Amazon S3 bucket or upload your file, refer to [Create your first S3 bucket](#) and [Upload an object to your bucket](#).

Java

```
import software.amazon.awssdk.auth.credentials.AwsCredentialsProvider;
import software.amazon.awssdk.auth.credentials.DefaultCredentialsProvider;
import software.amazon.awssdk.regions.Region;
import software.amazon.awssdk.services.transcribe.TranscribeClient;
import software.amazon.awssdk.services.transcribe.model.*;
import software.amazon.awssdk.services.transcribestreaming.model.LanguageCode;

public class TranscribeDemoApp {
    private static final Region REGION = Region.US_WEST_2;
    private static TranscribeClient client;

    public static void main(String args[]) {

        client = TranscribeClient.builder()
            .credentialsProvider(getCredentials())
            .region(REGION)
            .build();

        String transcriptionJobName = "my-first-transcription-job";
        String mediaType = "flac"; // can be other types
        Media myMedia = Media.builder()
            .mediaFileUri("s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-
file.flac")
            .build();

        String outputS3BucketName = "s3://DOC-EXAMPLE-BUCKET";
        // Create the transcription job request
        StartTranscriptionJobRequest request = StartTranscriptionJobRequest.builder()
            .transcriptionJobName(transcriptionJobName)
            .languageCode(LanguageCode.EN_US.toString())
            .mediaSampleRateHertz(16000)
            .mediaFormat(mediaType)
            .media(myMedia)
            .outputBucketName(outputS3BucketName)
            .build();

        // send the request to start the transcription job
        StartTranscriptionJobResponse startJobResponse =
client.startTranscriptionJob(request);

        System.out.println("Created the transcription job");
        System.out.println(startJobResponse.transcriptionJob());

        // Create the get job request
        GetTranscriptionJobRequest getJobRequest = GetTranscriptionJobRequest.builder()
            .transcriptionJobName(transcriptionJobName)
            .build();

        // send the request to get the transcription job including the job status
        GetTranscriptionJobResponse getJobResponse =
client.getTranscriptionJob(getJobRequest);

        System.out.println("Get the transcription job request");
        System.out.println(getJobResponse.transcriptionJob());
    }

    private static AwsCredentialsProvider getCredentials() {
        return DefaultCredentialsProvider.create();
    }
}
```

JavaScript

```
const { TranscribeClient, StartTranscriptionJobCommand } = require("@aws-sdk/client-transcribe"); // CommonJS import

const region = "us-west-2";
const credentials = {
  "accessKeyId": "",
  "secretAccessKey": "",
};

const input = {
  TranscriptionJobName: "my-first-transcription-job",
  LanguageCode: "en-US",
  Media: {
    MediaFileUri: "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
  },
  OutputBucketName: "DOC-EXAMPLE-BUCKET",
};

async function startTranscriptionRequest() {
  const transcribeConfig = {
    region,
    credentials
  };
  const transcribeClient = new TranscribeClient(transcribeConfig);
  const transcribeCommand = new StartTranscriptionJobCommand(input);
  try {
    const transcribeResponse = await transcribeClient.send(transcribeCommand);
    console.log("Transcription job created, the details:");
    console.log(transcribeResponse.TranscriptionJob);
  } catch(err) {
    console.log(err);
  }
}

startTranscriptionRequest();
```

Python

```
import time
import boto3

def transcribe_file(job_name, file_uri, transcribe_client):
    transcribe_client.start_transcription_job(
        TranscriptionJobName = job_name,
        Media = {
            'MediaFileUri': file_uri
        },
        MediaFormat = 'flac',
        LanguageCode = 'en-US'
    )

max_tries = 60
while max_tries > 0:
    max_tries -= 1
    job = transcribe_client.get_transcription_job(TranscriptionJobName = job_name)
    job_status = job['TranscriptionJob']['TranscriptionJobStatus']
    if job_status in ['COMPLETED', 'FAILED']:
        print(f"Job {job_name} is {job_status}.")
        if job_status == 'COMPLETED':
            print(
                f"Download the transcript from\n"
                f"\t{job['TranscriptionJob']['Transcript']['TranscriptFileUri']}")
```

```
        break
    else:
        print(f"Waiting for {job_name}. Current status is {job_status}.")
        time.sleep(10)

def main():
    transcribe_client = boto3.client('transcribe', region_name = 'us-west-2')
    file_uri = 's3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac'
    transcribe_file('Example-job', file_uri, transcribe_client)

if __name__ == '__main__':
    main()
```

Streaming transcriptions

You can create streaming transcriptions using a streamed media file or a live media stream.

Note that the standard AWS SDK for Python (Boto3) is not supported for Amazon Transcribe streaming. To start a streaming transcription using Python, use this [async Python SDK for Amazon Transcribe](#).

Java

The following example is a Java program that transcribes streaming audio.

To run this example, note the following:

- You must use the [AWS SDK for Java 2.x](#).
- Clients must use Java 1.8 to be compatible with the [AWS SDK for Java 2.x](#).
- The sample rate you specify must match the actual sample rate of your audio stream.

See also: [Retry client for Amazon Transcribe streaming \(Java SDK\)](#). This code manages the connection to Amazon Transcribe and retries sending data when there are errors on the connection. For example, if there is a transient error on the network, this client resends the request that failed.

```
public class TranscribeStreamingDemoApp {
    private static final Region REGION = Region.US_WEST_2;
    private static TranscribeStreamingAsyncClient client;

    public static void main(String args[]) throws URISyntaxException,
        ExecutionException, InterruptedException, LineUnavailableException {

        client = TranscribeStreamingAsyncClient.builder()
            .credentialsProvider(getCredentials())
            .region(REGION)
            .build();

        CompletableFuture<Void> result =
            client.startStreamTranscription(getRequest(16_000),
                new AudioStreamPublisher(getStreamFromMic()),
                getResponseHandler());

        result.get();
        client.close();
    }

    private static InputStream getStreamFromMic() throws LineUnavailableException {

        // Signed PCM AudioFormat with 16kHz, 16 bit sample size, mono
```

```
int sampleRate = 16000;
AudioFormat format = new AudioFormat(sampleRate, 16, 1, true, false);
DataLine.Info info = new DataLine.Info(TargetDataLine.class, format);

if (!AudioSystem.isLineSupported(info)) {
    System.out.println("Line not supported");
    System.exit(0);
}

TargetDataLine line = (TargetDataLine) AudioSystem.getLine(info);
line.open(format);
line.start();

InputStream audioStream = new AudioInputStream(line);
return audioStream;
}

private static AwsCredentialsProvider getCredentials() {
    return DefaultCredentialsProvider.create();
}

private static StartStreamTranscriptionRequest getRequest(Integer
mediaSampleRateHertz) {
    return StartStreamTranscriptionRequest.builder()
        .languageCode(LanguageCode.EN_US.toString())
        .mediaEncoding(MediaEncoding.PCM)
        .mediaSampleRateHertz(mediaSampleRateHertz)
        .build();
}

private static StartStreamTranscriptionResponseHandler getResponseHandler() {
    return StartStreamTranscriptionResponseHandler.builder()
        .onResponse(r -> {
            System.out.println("Received Initial response");
        })
        .onError(e -> {
            System.out.println(e.getMessage());
            StringWriter sw = new StringWriter();
            e.printStackTrace(new PrintWriter(sw));
            System.out.println("Error Occurred: " + sw.toString());
        })
        .onComplete(() -> {
            System.out.println("=== All records stream successfully ===");
        })
        .subscriber(event -> {
            List<Result> results = ((TranscriptEvent)
event).transcript().results();
            if (results.size() > 0) {
                if (!
results.get(0).alternatives().get(0).transcript().isEmpty()) {
                    System.out.println(results.get(0).alternatives().get(0).transcript());
                }
            }
        })
        .build();
}

private InputStream getStreamFromFile(String myMediaFileName) {
    try {
        File inputFile = new
File(getClass().getClassLoader().getResource(myMediaFileName).getFile());
        InputStream audioStream = new FileInputStream(inputFile);
        return audioStream;
    } catch (FileNotFoundException e) {
        throw new RuntimeException(e);
    }
}
```

```
    }  
}  
  
private static class AudioStreamPublisher implements Publisher<AudioStream> {  
    private final InputStream inputStream;  
    private static Subscription currentSubscription;  
  
    private AudioStreamPublisher(InputStream inputStream) {  
        this.inputStream = inputStream;  
    }  
  
    @Override  
    public void subscribe(Subscriber<? super AudioStream> s) {  
  
        if (this.currentSubscription == null) {  
            this.currentSubscription = new SubscriptionImpl(s, inputStream);  
        } else {  
            this.currentSubscription.cancel();  
            this.currentSubscription = new SubscriptionImpl(s, inputStream);  
        }  
        s.onSubscribe(currentSubscription);  
    }  
}  
  
public static class SubscriptionImpl implements Subscription {  
    private static final int CHUNK_SIZE_IN_BYTES = 1024 * 1;  
    private final Subscriber<? super AudioStream> subscriber;  
    private final InputStream inputStream;  
    private ExecutorService executor = Executors.newFixedThreadPool(1);  
    private AtomicLong demand = new AtomicLong(0);  
  
    SubscriptionImpl(Subscriber<? super AudioStream> s, InputStream inputStream) {  
        this.subscriber = s;  
        this.inputStream = inputStream;  
    }  
  
    @Override  
    public void request(long n) {  
        if (n <= 0) {  
            subscriber.onError(new IllegalArgumentException("Demand must be  
positive"));  
        }  
  
        demand.getAndAdd(n);  
  
        executor.submit(() -> {  
            try {  
                do {  
                    ByteBuffer audioBuffer = getNextEvent();  
                    if (audioBuffer.remaining() > 0) {  
                        AudioEvent audioEvent = audioEventFromBuffer(audioBuffer);  
                        subscriber.onNext(audioEvent);  
                    } else {  
                        subscriber.onComplete();  
                        break;  
                    }  
                } while (demand.decrementAndGet() > 0);  
            } catch (Exception e) {  
                subscriber.onError(e);  
            }  
        });  
    }  
  
    @Override  
    public void cancel() {
```

```
        executor.shutdown();
    }

    private ByteBuffer getNextEvent() {
        ByteBuffer audioBuffer = null;
        byte[] audioBytes = new byte[CHUNK_SIZE_IN_BYTES];

        int len = 0;
        try {
            len = inputStream.read(audioBytes);

            if (len <= 0) {
                audioBuffer = ByteBuffer.allocate(0);
            } else {
                audioBuffer = ByteBuffer.wrap(audioBytes, 0, len);
            }
        } catch (IOException e) {
            throw new UncheckedIOException(e);
        }

        return audioBuffer;
    }

    private AudioEvent audioEventFromBuffer(ByteBuffer bb) {
        return AudioEvent.builder()
            .audioChunk(SdkBytes.fromByteBuffer(bb))
            .build();
    }
}
}
```

JavaScript

```
const {
    TranscribeStreamingClient,
    StartStreamTranscriptionCommand,
} = require("@aws-sdk/client-transcribe-streaming");
const { createReadStream } = require("fs");
const { join } = require("path");

const audio = createReadStream(join(__dirname, "my-media-file.flac"), { highWaterMark:
    1024 * 16});

const LanguageCode = "en-US";
const MediaEncoding = "pcm";
const MediaSampleRateHertz = "16000";
const credentials = {
    "accessKeyId": "",
    "secretAccessKey": "",
};

async function startRequest() {
    const client = new TranscribeStreamingClient({
        region: "us-west-2",
        credentials
    });

    const params = {
        LanguageCode,
        MediaEncoding,
        MediaSampleRateHertz,
        AudioStream: (async function* () {
            for await (const chunk of audio) {
                yield {AudioEvent: {AudioChunk: chunk}};
            }
        })(),
    };
}
```

```
};
const command = new StartStreamTranscriptionCommand(params);
// Send transcription request
const response = await client.send(command);
// Start to print response
try {
  for await (const event of response.TranscriptResultStream) {
    console.log(JSON.stringify(event));
  }
} catch(err) {
  console.log("error")
  console.log(err)
}
}
startRequest();
```

Python

The following example is a Python program that transcribes streaming audio.

To run this example, note the following:

- You must use this [SDK for Python](#) .
- The sample rate you specify must match the actual sample rate of your audio stream.

```
import asyncio
# This example uses aiofile for asynchronous file reads.
# It's not a dependency of the project but can be installed
# with `pip install aiofile`.
import aiofile

from amazon_transcribe.client import TranscribeStreamingClient
from amazon_transcribe.handlers import TranscriptResultStreamHandler
from amazon_transcribe.model import TranscriptEvent

"""
Here's an example of a custom event handler you can extend to
process the returned transcription results as needed. This
handler will simply print the text out to your interpreter.
"""
class MyEventHandler(TranscriptResultStreamHandler):
    async def handle_transcript_event(self, transcript_event: TranscriptEvent):
        # This handler can be implemented to handle transcriptions as needed.
        # Here's an example to get started.
        results = transcript_event.transcript.results
        for result in results:
            for alt in result.alternatives:
                print(alt.transcript)

async def basic_transcribe():
    # Set up our client with your chosen Region
    client = TranscribeStreamingClient(region = "us-west-2")

    # Start transcription to generate async stream
    stream = await client.start_stream_transcription(
        language_code = "en-US",
        media_sample_rate_hz = 16000,
        media_encoding = "pcm",
    )

    async def write_chunks():
        # NOTE: For pre-recorded files longer than 5 minutes, the sent audio
```

```
# chunks should be rate limited to match the real-time bitrate of the
# audio stream to avoid signing issues.
async with aiofile.AIOFile('filepath/my-media-file.flac', 'rb') as afp:
    reader = aiofile.Reader(afp, chunk_size = 1024 * 16)
    async for chunk in reader:
        await stream.input_stream.send_audio_event(audio_chunk = chunk)
    await stream.input_stream.end_stream()

# Instantiate our handler and start processing events
handler = MyEventHandler(stream.output_stream)
await asyncio.gather(write_chunks(), handler.handle_events())

loop = asyncio.get_event_loop()
loop.run_until_complete(basic_transcribe())
loop.close()
```

Using this service with an AWS SDK

AWS software development kits (SDKs) are available for many popular programming languages. Each SDK provides an API, code examples, and documentation that make it easier for developers to build applications in their preferred language.

SDK documentation	Code examples
AWS SDK for C++	AWS SDK for C++ code examples
AWS SDK for Go	AWS SDK for Go code examples
AWS SDK for Java	AWS SDK for Java code examples
AWS SDK for JavaScript	AWS SDK for JavaScript code examples
AWS SDK for .NET	AWS SDK for .NET code examples
AWS SDK for PHP	AWS SDK for PHP code examples
AWS SDK for Python (Boto3)	AWS SDK for Python (Boto3) code examples
AWS SDK for Ruby	AWS SDK for Ruby code examples

For examples specific to this service, see [Code examples for Amazon Transcribe using AWS SDKs](#) (p. 202).

Example availability

Can't find what you need? Request a code example by using the **Provide feedback** link at the bottom of this page.

Transcribing with HTTP or WebSockets

Amazon Transcribe supports HTTP for both batch (HTTP/1.1) and streaming (HTTP/2) transcriptions. WebSockets are supported for streaming transcriptions.

Both HTTP and WebSockets require you to authenticate your request using AWS Signature Version 4 headers. Refer to [Signing AWS API requests](#) for more information.

You can make a batch HTTP request using the following headers:

- host
- x-amz-target
- content-type
- x-amz-content-sha256
- x-amz-date
- authorization

Here's an example of a `StartTranscriptionJob` request:

```
POST /transcribe HTTP/1.1
host: transcribe.us-west-2.amazonaws.com
x-amz-target: com.amazonaws.transcribe.Transcribe.StartTranscriptionJob
content-type: application/x-amz-json-1.1
x-amz-content-sha256: string
x-amz-date: YYYYMMDDTHHMMSSZ
authorization: AWS4-HMAC-SHA256 Credential=access-key/YYYYMMSS/us-west-2/transcribe/
aws4_request, SignedHeaders=content-type;host;x-amz-content-sha256;x-amz-date;x-amz-
target;x-amz-security-token, Signature=string

{
  "TranscriptionJobName": "my-first-transcription-job",
  "LanguageCode": "en-US",
  "Media": {
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
  },
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",
  "OutputKey": "my-output-files/"
}
```

Additional operations and parameters are listed in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section. Other signature elements are detailed in [Elements of an AWS Signature Version 4 request](#).

Streaming transcriptions using HTTP/2 and WebSockets are more complicated. Detailed instructions are provided in these sections:

- [Setting up an HTTP/2 stream \(p. 82\)](#)
- [Setting up a WebSocket stream \(p. 86\)](#)

We recommend reviewing the [Transcribing streaming audio \(p. 76\)](#) chapter before setting up your first HTTP/2 or WebSocket stream.

Transcribing streaming audio

Using Amazon Transcribe streaming, you can produce real-time transcriptions for your media content. Unlike batch transcriptions, which involve uploading media files, streaming media is delivered to Amazon Transcribe in real time. Amazon Transcribe then returns a transcript, also in real time.

Streaming can include pre-recorded media (movies, music, and podcasts) and real-time media (live news broadcasts). Common streaming use cases for Amazon Transcribe include live closed captioning for sporting events and real-time monitoring of call center audio.

Streaming content is delivered as a series of sequential data packets, or 'chunks,' that Amazon Transcribe transcribes instantaneously. The advantages to using streaming over batch include real-time speech-to-text capabilities in your applications and faster transcription times. However, this increased speed may have accuracy limitations in some cases.

Amazon Transcribe offers the following options for streaming:

- [AWS SDKs](#)
- [HTTP/2 \(p. 82\)](#)
- [WebSockets \(p. 86\)](#)
- [AWS Management Console](#)

We recommend using the AWS Management Console or an SDK, rather than using HTTP/2 or WebSockets directly.

Audio formats supported for streaming transcriptions are:

- FLAC
- OPUS-encoded audio in an Ogg container
- PCM (only signed 16-bit little-endian audio formats, which does **not** include WAV)

Lossless formats (FLAC or PCM) are recommended.

Note

Streaming transcriptions are not supported with all languages. See [Supported languages and language-specific features \(p. 3\)](#) for details.

Streaming and partial results

Because streaming works in real time, transcripts are produced in *partial results*. Amazon Transcribe breaks up the incoming audio stream based on natural speech segments, such as a change in speaker or a pause in the audio. The transcription is returned to your application in a stream of transcription events, with each response containing more transcribed speech until an entire segment is transcribed.

An approximation of this is shown in the following code block. You can view this process in action by signing into the [AWS Management Console](#), selecting **Real-time transcription**, and speaking into your microphone. Watch the **Transcription output** pane as you speak.

In this example, each line is the partial result of an audio segment.

```
Welcome.  
Welcome to.
```

```
Welcome to Amazon.  
Welcome to Amazon transcribe.  
Welcome to Amazon transcribe.
```

These partial results are present in your transcription output within the [Results](#) objects. Also in this object block is an **IsPartial** field. If this field is true, your transcription segment is not yet complete. You can view the difference between an incomplete and a complete segment below:

```
"IsPartial": true (incomplete segment)  
  
"Transcript": "Welcome to Amazon."  
  
"EndTime": 2.695,  
"IsPartial": true,  
"ResultId": "55421b61-7cc6-4ef1-8d55-f1c128f01d30",  
"StartTime": 1.18  
  
"IsPartial": false (complete segment)  
  
"Transcript": "Welcome to Amazon transcribe."  
  
"EndTime": 4.775,  
"IsPartial": false,  
"ResultId": "55421b61-7cc6-4ef1-8d55-08a97159870e",  
"StartTime": 1.18
```

Each word within a complete segment has an associated confidence score, which is a value between 0 and 1. A larger value indicates a greater likelihood that the word is correctly transcribed.

Tip

The `StartTime` and `EndTime` of an audio segment can be used to synchronize transcription output with video dialogue.

If you're running an application that requires low latency, you may want to use [partial-result stabilization](#) (p. 77).

Partial-result stabilization

Amazon Transcribe starts returning transcription results as soon as you start streaming your audio. It returns these partial results incrementally until it generates a finished result at the level of a natural speech segment. A natural speech segment is continuous speech that contains a pause or a change in speaker.

Amazon Transcribe continues outputting partial results until it generates the final transcription result for a speech segment. Because speech recognition may revise words as it gains more context, streaming transcriptions can change slightly with each new partial result output.

This process gives you two options for each speech segment:

- Wait for the finished segment
- Use the segment's partial results

Partial result stabilization changes how Amazon Transcribe produces the final transcription result for each complete segment. When activated, only the last few words from the partial results can change. Because of this, transcription accuracy may be affected. However, your transcript is returned faster than without partial-results stabilization. This reduction in latency may be beneficial when subtitling videos or generating captions for live streams.

The following examples show how the same audio stream is handled when partial-results stabilization is not activated and when it is. Note that you can set the stability level to low, medium, or high. Low stability provides the highest accuracy. High stability transcribes faster, but with slightly lower accuracy.

"Transcript":	"EndTime":	"IsPartial":
Partial-result stabilization not enabled		
Welcome. Welcome to. Welcome to Amazon. Welcome to Amazon. transcribe. Welcome to Amazon transcribe. Welcome to Amazon transcribe. Welcome to Amazon transcribe. Welcome to Amazon transcribe.	0.515 1.015 1.515 2.015 2.515 3.015 3.515 3.595 3.595	true true true true true true true true true false
Partial-result stabilization enabled (high stability)		
Welcome. Welcome to Welcome to Amas. Welcome to Amazon. Welcome to Amazon. Welcome to Amazon transcribe. Welcome to Amazon transcribe.	0.515 1.015 1.515 2.015 2.515 2.865 2.865	true true true true true true true false

When you activate partial-result stabilization, Amazon Transcribe uses a `stable` field to indicate whether an item is stable, where 'item' refers to a transcribed word or phrase. Values for `stable` are either `true`, indicating that the item is stable, or `false`, indicating that the item is not stable. Items flagged as `false` (not stable) are more likely to change as your segment is transcribed. Conversely, items flagged as `true` (stable) don't change.

You can choose to render non-stable words so your captions align with speech. Even if captions change slightly as context is added, this is a better user experience than periodic text bursts, which may or may not align with speech.

You can also choose to display non-stable words in a different format, such as italics, to indicate to viewers that these words may change. Displaying partial results limits the amount of text displayed at a given time. This can be important when you're dealing with space constraints, as with video captions.

Dive deeper with the AWS Machine Learning Blog

To learn more about improving accuracy with real-time transcriptions, see:

- [Improve the streaming transcription experience with Amazon Transcribe partial results stabilization](#)
- ["What was that?" Increasing subtitle accuracy for live broadcasts using Amazon Transcribe](#)

Partial-result stabilization example output

The following example output shows `Stable` flags for an incomplete segment (`"IsPartial": true`). You can see that the words `"to Amazon"` are not stable and therefore could change before the segment is finalized.

```
"Transcript": {
  "Results": [
    {
      "Alternatives": [
        {
          "Items": [
            {
              "Content": "Welcome",
              "EndTime": 2.4225,
              "Stable": true,
              "StartTime": 1.65,
              "Type": "pronunciation",
              "VocabularyFilterMatch": false
            },
            {
              "Content": "to",
              "EndTime": 2.8325,
              "Stable": false,
              "StartTime": 2.4225,
              "Type": "pronunciation",
              "VocabularyFilterMatch": false
            },
            {
              "Content": "Amazon",
              "EndTime": 3.635,
              "Stable": false,
              "StartTime": 2.8325,
              "Type": "pronunciation",
              "VocabularyFilterMatch": false
            },
            {
              "Content": ".",
              "EndTime": 3.635,
              "Stable": false,
              "StartTime": 3.635,
              "Type": "punctuation",
              "VocabularyFilterMatch": false
            }
          ],
          "Transcript": "Welcome to Amazon."
        }
      ],
      "EndTime": 4.165,
      "IsPartial": true,
      "ResultId": "364f1c74-5e21-4333-af74-b6531a6b9158",
      "StartTime": 1.65
    }
  ]
}
```

Setting up a streaming transcription

This section expands on the main [streaming \(p. 76\)](#) section and is intended to provide information for users who want to set up their stream with HTTP/2 or WebSockets directly, rather than with an AWS SDK. The information in this section can also be used to build your own SDK.

To start streaming using an AWS SDK, see [Transcribing with the AWS SDKs \(p. 66\)](#).

Event stream encoding

Amazon Transcribe uses a format called event stream encoding for streaming transcriptions.

Event stream encoding provides bidirectional communication between a client and a server. Data frames sent to the Amazon Transcribe streaming service are encoded in this format. The response from Amazon Transcribe also uses this encoding.

Each message consists of two sections: the prelude and the data. The prelude consists of:

1. The total byte length of the message
2. The combined byte length of all headers

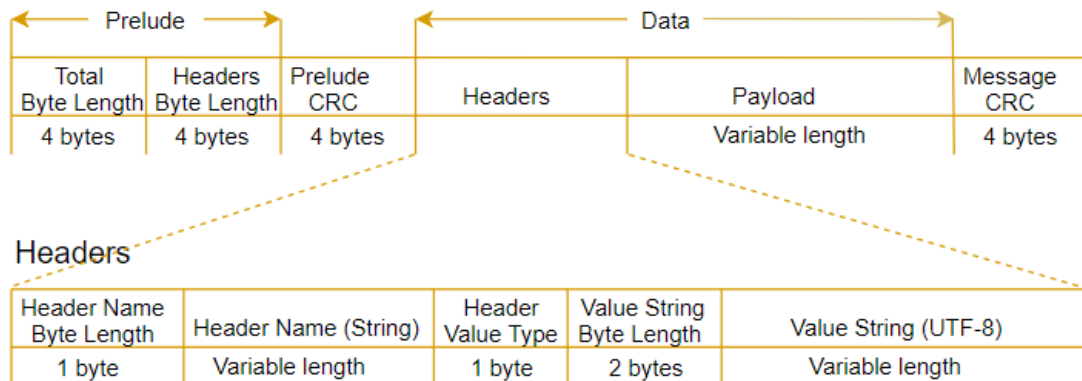
The data section consists of:

1. Headers
2. Payload

Each section ends with a 4-byte big-endian integer cyclic redundancy check (CRC) checksum. The message CRC checksum is for both the prelude section and the data section. Amazon Transcribe uses CRC32 (often referred to as GZIP CRC32) to calculate both CRCs. For more information about CRC32, see [GZIP file format specification version 4.3](#).

Total message overhead, including the prelude and both checksums, is 16 bytes.

The following diagram shows the components that make up a message and a header. There are multiple headers per message.



Each message contains the following components:

- **Prelude:** Consists of two, 4-byte fields, for a fixed total of 8 bytes.
 - *First 4 bytes:* The big-endian integer byte-length of the entire message, inclusive of this 4-byte length field.

- *Second 4 bytes*: The big-endian integer byte-length of the 'headers' portion of the message, excluding the 'headers' length field itself.
- **Prelude CRC**: The 4-byte CRC checksum for the prelude portion of the message, excluding the CRC itself. The prelude has a separate CRC from the message CRC to ensure that Amazon Transcribe can detect corrupted byte-length information immediately without causing errors, such as buffer overruns.
- **Headers**: Metadata annotating the message; for example, message type and content type. Messages have multiple headers, which are key:value pairs, where the key is a UTF-8 string. Headers can appear in any order in the 'headers' portion of the message, and each header can appear only once.
- **Payload**: The audio content to be transcribed.
- **Message CRC**: The 4-byte CRC checksum from the start of the message to the start of the checksum. That is, everything in the message except the CRC itself.

The header frame is the authorization frame for the streaming transcription. Amazon Transcribe uses the value of the authorization header as the seed for generating a chain of authorization headers for the data frames in the request.

Each header contains the following components; there are multiple headers per frame.

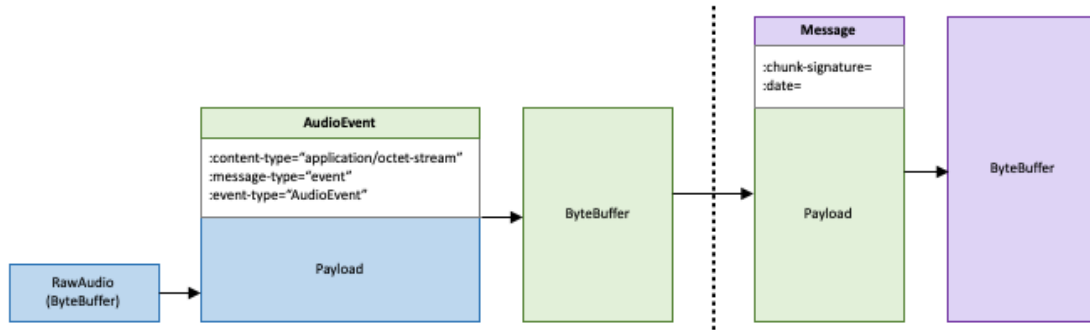
- **Header name byte-length**: The byte-length of the header name.
- **Header name**: The name of the header that indicates the header type. For valid values, see the following frame descriptions.
- **Header value type**: A number indicating the header value. The following list shows the possible values for the header and what they indicate.
 - 0 – TRUE
 - 1 – FALSE
 - 2 – BYTE
 - 3 – SHORT
 - 4 – INTEGER
 - 5 – LONG
 - 6 – BYTE ARRAY
 - 7 – STRING
 - 8 – TIMESTAMP
 - 9 – UUID
- **Value string byte length**: The byte length of the header value string.
- **Header value**: The value of the header string. Valid values for this field depend on the type of header. See [Setting up an HTTP/2 stream \(p. 82\)](#) or [Setting up a WebSocket stream \(p. 86\)](#) for more information.

Data frames

Each streaming request contains one or more data frames. There are two steps to creating a data frame:

1. Combine raw audio data with metadata to create the payload of your request.
2. Combine the payload with a signature to form the event message that is sent to Amazon Transcribe.

The following diagram shows how this works.



You can find more information on data frames in the [Setting up an HTTP/2 stream \(p. 82\)](#) and [Setting up a WebSocket stream \(p. 86\)](#) sections.

Setting up an HTTP/2 stream

The key components for an [HTTP/2 protocol](#) for streaming transcription requests with Amazon Transcribe are:

- A header frame. This contains the HTTP/2 headers for your request, and a signature in the authorization header that Amazon Transcribe uses as a seed signature to sign the data frames.
- One or more message frames in [event stream encoding \(p. 80\)](#) that contain metadata and raw audio bytes.
- An end frame. This is a signed message in [event stream encoding \(p. 80\)](#) with an empty body.

Note

Amazon Transcribe only supports one stream per HTTP/2 session. If you attempt to use multiple streams, your transcription request fails.

1. Attach the following policy to the IAM role that makes the request. See [Adding IAM policies](#) for more information.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "my-transcribe-http2-policy",
      "Effect": "Allow",
      "Action": "transcribe:StartStreamTranscription",
      "Resource": "*"
    }
  ]
}
```

2. To start the session, send an HTTP/2 request to Amazon Transcribe.

```
POST /stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
X-Amz-Target: com.amazonaws.transcribe.Transcribe.StartStreamTranscription
Content-Type: application/vnd.amazon.eventstream
X-Amz-Content-Sha256: string
X-Amz-Date: YYYYMMDDTHHMMSSZ
Authorization: AWS4-HMAC-SHA256 Credential=access-key/YYYYMMDD/us-west-2/transcribe/
aws4_request, SignedHeaders=content-type;host;x-amz-content-sha256;x-amz-date;x-amz-
target;x-amz-security-token, Signature=string
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
```



```
x-amzn-transcribe-sample-rate: 16000
transfer-encoding: chunked
```

Additional operations and parameters are listed in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

Amazon Transcribe sends the following response:

```
HTTP/2.0 200
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
x-amzn-request-id: 8a08df7d-5998-48bf-a303-484355b4ab4e
x-amzn-transcribe-session-id: b4526fcf-5eee-4361-8192-d1cb9e9d6887
content-type: application/json
```

3. Create an audio event that contains your audio data. Combine the headers—described in the following table—with a chunk of audio bytes in an event-encoded message. To create the payload for the event message, use a buffer in raw-byte format.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	24	application/octet-stream
11	:event-type	7	10	AudioEvent
13	:message-type	7	5	event

Binary data in this example request are base64-encoded. In an actual request, data are raw bytes.

```
:content-type: "application/vnd.amazon.eventstream"
:event-type: "AudioEvent"
:message-type: "event"
UklGRjzxPQBXQVZFZm10IBAAAAABAAEAgD4AAAB9AAACABAAZGF0YVVTwPQAAAAAAAAAAAAAAAAAD//wIA/f8EAA==
```

4. Create an audio message that contains your audio data.
 - a. Your audio message data frame contains event-encoding headers that include the current date and a signature for the audio chunk and the audio event.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value
16	:chunk-signature	6	varies	generated signature
5	:date	8	8	timestamp

Binary data in this request are base64-encoded. In an actual request, data are raw bytes.

```
:date: 2019-01-29T01:56:17.291Z
:chunk-signature: signature
```

```
AAAA0gAAAIKVoRFcTtcjb250ZW50LXR5cGUHABhbcHBsaWNhdGlvbi9vY3RldC1zdHJlYW0LOmV2ZW50LXR5cGUHAApBdWRpb0V2ZW50DTptZXXNzYWdlLXR5cGUHAAV1dmVudAxDb256ZW50LVR5cGUHABphcHBsaWNhdGlvbi94LWFteil1qc29uLTFuMVJJKyY88T0AV0FWRWZtdCAQAAAAAQAABAIA  
+AAAAfQAAAQAQRhdGFU8D0AAAAA  
AAAAAAAAAAAA//8CAP3/BAC7QLFf
```

- b. Construct a string to sign, as outlined in [Create a string to sign for Signature Version 4](#). Your string follows this format:

```
String stringToSign =  
"AWS4-HMAC-SHA256" +  
"\n" +  
DateTime +  
"\n" +  
Keypath +  
"\n" +  
Hex(priorSignature) +  
"\n" +  
HexHash(nonSignatureHeaders) +  
"\n" +  
HexHash(payload);
```

- **DateTime:** The date and time the signature is created. The format is YYYYMMDDTHHMMSSZ, where YYYY=year, MM=month, DD=day, HH=hour, MM=minute, SS=seconds, and 'T' and 'Z' are fixed characters. For more information, refer to [Handling Dates in Signature Version 4](#).
 - **Keypath:** The signature scope in the format date/region/service/aws4_request. For example, 20220127/us-west-2/transcribe/aws4_request.
 - **Hex:** A function that encodes input into a hexadecimal representation.
 - **priorSignature:** The signature for the previous frame. For the first data frame, use the signature of the header frame.
 - **HexHash:** A function that first creates a SHA-256 hash of its input and then uses the Hex function to encode the hash.
 - **nonSignatureHeaders:** The DateTime header encoded as a string.
 - **payload:** The byte buffer containing the audio event data.
- c. Derive a signing key from your AWS secret access key and use it to sign the `stringToSign`. For a greater degree of protection, the derived key is specific to the date, service, and AWS Region. For more information, see [Calculate the signature for AWSSignature Version 4](#).

Make sure you implement the `GetSignatureKey` function to derive your signing key. If you have not yet derived a signing key, refer to [Examples of how to derive a signing key for Signature Version 4](#).

```
String signature = HMACSHA256(derivedSigningKey, stringToSign);
```

- **HMACSHA256:** A function that creates a signature using the SHA-256 hash function.
- **derivedSigningKey:** The Signature Version 4 signing key.
- **stringToSign:** The string you calculated for the data frame.

After you've calculated the signature for the data frame, construct a byte buffer containing the date, signature, and audio event payload. Send the byte array to Amazon Transcribe for transcription.

5. To indicate the audio stream is complete, send an end frame (an empty data frame) that contains only the date and signature. You construct this end frame the same way you construct a data frame.

Amazon Transcribe responds with a stream of transcription events, sent to your application. This response is event stream encoded. It contains the standard prelude and the following headers.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	16	application/json
11	:event-type	7	15	TranscriptEvent
13	:message-type	7	5	event

The events are sent in raw-byte format. In this example, the bytes are base64-encoded.

```
AAAAUwAAAEp1RHpYBTpkYXR1C AAAAWiXUkMLEDPjaHVuay1zaWduYXR1cmUGACct6Zy+uymwEK2SrLp/zVBI5eGn83jdBwCaRUBJA+eaDafqjqI=
```

To see the transcription results, decode the raw bytes using event stream encoding.

```
:content-type: "application/vnd.amazon.eventstream"
:event-type: "TranscriptEvent"
:message-type: "event"

{
  "Transcript":
  {
    "Results":
    [
      results
    ]
  }
}
```

- To end your stream, send an empty audio event to Amazon Transcribe. Create the audio event exactly like any other, except with an empty payload. Sign the event and include the signature in the :chunk-signature header, as follows:

```
:date: 2019-01-29T01:56:17.291Z
:chunk-signature: signature
```

Handling HTTP/2 streaming errors

If an error occurs when processing your media stream, Amazon Transcribe sends an exception response. The response is event stream encoded.

The response contains the standard prelude and the following headers:

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	16	application/json
11	:event-type	7	19	BadRequestException

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:message-type	7	9	exception

When the exception response is decoded, it contains the following information:

```
:content-type: "application/vnd.amazon.eventstream"  
:event-type: "BadRequestException"  
:message-type: "exception"
```

Exception message

Setting up a WebSocket stream

The key components for a [WebSocket protocol](#) for streaming transcription requests with Amazon Transcribe are:

- The upgrade request. This contains the query parameters for your request, and a signature that Amazon Transcribe uses as a seed signature to sign the data frames.
- One or more message frames in [event stream encoding \(p. 80\)](#) that contain metadata and raw audio bytes.
- An end frame. This is a signed message in [event stream encoding \(p. 80\)](#) with an empty body.

Note

Amazon Transcribe only supports one stream per WebSocket session. If you attempt to use multiple streams, your transcription request fails.

1. Attach the following policy to the IAM role that makes the request. See [Adding IAM policies](#) for more information.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Sid": "my-transcribe-websocket-policy",  
      "Effect": "Allow",  
      "Action": "transcribe:StartStreamTranscriptionWebSocket",  
      "Resource": "*"  
    }  
  ]  
}
```

2. To start the session, create a pre-signed URL in the following format. Line breaks have been added for readability.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/stream-transcription-  
websocket?  
&X-Amz-Algorithm=AWS4-HMAC-SHA256  
&X-Amz-Credential=access-key%2FYYYYMMDD%2Fus-west-2%2Ftranscribe%2Faws4_request  
&X-Amz-Date=YYYYMMDDTHHMMSSZ  
&X-Amz-Expires=250  
&X-Amz-Security-Token=security-token  
&X-Amz-Signature=string  
&X-Amz-SignedHeaders=content-type%3Bhost%3Bx-amz-date  
&language-code=en-US  
&media-encoding=flac
```

```
&sample-rate=16000
```

Additional operations and parameters are listed in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

To construct the URL for your request and create the [Signature Version 4 signature](#), refer to the following steps. Examples are in pseudocode.

- a. Create a canonical request. A canonical request is a string that includes information from your request in a standardized format. This ensures that when AWS receives the request, it can calculate the same signature you created for your URL. For more information, see [Create a Canonical Request for Signature Version 4](#).

```
# HTTP verb
method = "GET"
# Service name
service = "transcribe"
# Region
region = "us-west-2"
# Amazon Transcribe streaming endpoint
endpoint = "wss://transcribestreaming.us-west-2.amazonaws.com:8443"
# Host
host = "transcribestreaming.us-west-2.amazonaws.com:8443"
# Date and time of request
amz-date = YYYYMMDDTHHMMSSZ
# Date without time for credential scope
datestamp = YYYYMMDD
```

- b. Create a canonical URI, which is the part of the URI between the domain and the query string.

```
canonical_uri = "/stream-transcription-websocket"
```

- c. Create the canonical headers and signed headers. Note the trailing `\n` in the canonical headers.
 - Append the lowercase header name followed by a colon (`:`).
 - Append a comma-separated list of values for that header. Do not sort values in headers that have multiple values.
 - Append a new line (`\n`).

```
canonical_headers = "host:" + host + "\n"
signed_headers = "host"
```

- d. Match the algorithm to the hashing algorithm. You must use SHA-256.

```
algorithm = "AWS4-HMAC-SHA256"
```

- e. Create the credential scope, which scopes the derived key to the date, AWS Region, and service. For example, `20220127/us-west-2/transcribe/aws4_request`.

```
credential_scope = datestamp + "/" + region + "/" + service + "/" + "aws4_request"
```

- f. Create the canonical query string. Query string values must be URI-encoded and sorted by name.
 - Sort the parameter names by character code point in ascending order. Parameters with duplicate names should be sorted by value. For example, a parameter name that begins with the uppercase letter F precedes a parameter name that begins with the lowercase letter b.

- Do not URI-encode any of the unreserved characters that [RFC 3986](#) defines: A-Z, a-z, 0-9, hyphen (-), underscore (_), period (.), and tilde (~).
- Percent-encode all other characters with %XY, where X and Y are hexadecimal characters (0-9 and uppercase A-F). For example, the space character must be encoded as %20 (don't include '+', as some encoding schemes do); extended UTF-8 characters must be in the form %XY%ZA%BC.
- Double-encode any equals (=) characters in parameter values.

```
canonical_querystring = "X-Amz-Algorithm=" + algorithm
canonical_querystring += "&X-Amz-Credential=" + URI-encode(access key + "/" +
  credential_scope)
canonical_querystring += "&X-Amz-Date=" + amz_date
canonical_querystring += "&X-Amz-Expires=250"
canonical_querystring += "&X-Amz-Security-Token=" + token
canonical_querystring += "&X-Amz-SignedHeaders=" + signed_headers
canonical_querystring += "&language-code=en-US&media-encoding=flac&sample-
  rate=16000"
```

- g. Create a hash of the payload. For a GET request, the payload is an empty string.

```
payload_hash = HashSHA256("").Encode("utf-8").HexDigest()
```

- h. Combine the following elements to create the canonical request.

```
canonical_request = method + '\n'
  + canonical_uri + '\n'
  + canonical_querystring + '\n'
  + canonical_headers + '\n'
  + signed_headers + '\n'
  + payload_hash
```

3. Create the string to sign, which contains meta information about your request. You use the string to sign in the next step when you calculate the request signature. For more information, see [Create a String to Sign for Signature Version 4](#).

```
string_to_sign=algorithm + "\n"
  + amz_date + "\n"
  + credential_scope + "\n"
  + HashSHA256(canonical_request.Encode("utf-8")).HexDigest()
```

4. Calculate the signature. To do this, derive a signing key from your AWS secret access key. For a greater degree of protection, the derived key is specific to the date, service, and AWS Region. Use this derived key to sign the request. For more information, see [Calculate the Signature for AWS Signature Version 4](#).

Make sure you implement the `GetSignatureKey` function to derive your signing key. If you have not yet derived a signing key, refer to [Examples of how to derive a signing key for Signature Version 4](#).

```
#Create the signing key
signing_key = GetSignatureKey(secret_key, timestamp, region, service)

# Sign the string_to_sign using the signing key
signature = HMAC.new(signing_key, (string_to_sign).Encode("utf-8"), Sha256()).HexDigest
```

The function `HMAC(key, data)` represents an HMAC-SHA256 function that returns results in binary format.

5. Add signing information to the request and create the request URL.

After you calculate the signature, add it to the query string. For more information, see [Add the Signature to the Request](#).

First, add the authentication information to the query string.

```
canonical_querystring += "&X-Amz-Signature=" + signature
```

Second, create the URL for the request.

```
request_url = endpoint + canonical_uri + "?" + canonical_querystring
```

Use the request URL with your WebSocket library to make the request to Amazon Transcribe.

6. The request to Amazon Transcribe must include the following headers. Typically these headers are managed by your WebSocket client library.

```
Host: transcribestreaming.us-west-2.amazonaws.com:8443
Connection: Upgrade
Upgrade: websocket
Origin: URI-of-WebSocket-client
Sec-WebSocket-Version: 13
Sec-WebSocket-Key: randomly-generated-string
```

7. When Amazon Transcribe receives your WebSocket request, it responds with a WebSocket upgrade response. Typically your WebSocket library manages this response and sets up a socket for communications with Amazon Transcribe.

The following is the response from Amazon Transcribe. Line breaks have been added for readability.

```
HTTP/1.1 101 WebSocket Protocol Handshake

Connection: upgrade
Upgrade: websocket
websocket-origin: wss://transcribestreaming.us-west-2.amazonaws.com:8443
websocket-location: transcribestreaming.us-west-2.amazonaws.com:8443/stream-
transcription-websocket?
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe
%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Signature=Signature Version 4 signature
&X-Amz-SignedHeaders=host
&language-code=en-US
&session-id=String
&media-encoding=flac
&sample-rate=16000
x-amzn-RequestId: RequestId
Strict-Transport-Security: max-age=31536000
sec-websocket-accept: hash-of-the-Sec-WebSocket-Key-header
```

8. Make your WebSocket streaming request.

After the WebSocket connection is established, the client can start sending a sequence of audio frames, each encoded using [event stream encoding \(p. 80\)](#).

Each data frame contains three headers combined with a chunk of raw audio bytes; the following table describes these headers.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	24	application/octet-stream
11	:event-type	7	10	AudioEvent
13	:message-type	7	5	event

9. To end the data stream, send an empty audio chunk in an event stream encoded message.

The response contains event stream encoded raw bytes in the payload. It contains the standard prelude and the following headers.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	16	application/json
11	:event-type	7	15	TranscriptEvent
13	:message-type	7	5	event

When you decode the binary response, you end up with a JSON structure containing the transcription results.

Handling WebSocket streaming errors

If an exception occurs while processing your request, Amazon Transcribe responds with a terminal WebSocket frame containing an event stream encoded response. This response contains the headers described in the following table; the body of the response contains a descriptive error message. After sending the exception response, Amazon Transcribe sends a close frame.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	16	application/json
15	:exception-type	7	varies	varies, see below
13	:message-type	7	9	exception

The `exception-type` header contains one of the following values:

- **BadRequestException:** There was a client error when the stream was created, or an error occurred while streaming data. Make sure your client is ready to accept data and try your request again.
- **InternalFailureException:** Amazon Transcribe had a problem during the handshake with the client. Try your request again.
- **LimitExceededException:** The client exceeded the concurrent stream limit. For more information, see [Amazon Transcribe Limits](#). Reduce the number of streams you're transcribing.
- **UnrecognizedClientException:** The WebSocket upgrade request was signed with an incorrect access key or secret key. Make sure you're correctly creating the access key and try your request again.

Amazon Transcribe can also return any of the common service errors. For a list, see [Common Errors](#).

Job queueing

Job queueing enables you to submit more transcription job requests than can be concurrently processed. Without job queueing, once you reach the quota of allowed concurrent requests, you must wait until one or more requests are completed before submitting a new request.

Job queueing is optional for transcription job requests. Call Analytics job requests have job queueing enabled automatically.

If you enable job queueing, Amazon Transcribe creates a queue that contains all requests that exceed your limit. As soon as a request is completed, a new request is pulled from your queue and processed. Queued requests are processed in a FIFO (first in, first out) order.

You can add up to 10,000 jobs to your queue. If you exceed this limit, you get a `LimitExceededConcurrentJobException` error. To maintain optimal performance, Amazon Transcribe only uses up to 90 percent of your quota (a bandwidth ratio of 0.9) to process queued jobs. Note that this is a default value that can be increased.

Tip

You can find a list of default limits and quotas for Amazon Transcribe resources in the [Quotas \(p. 42\)](#) section. Some of these defaults can be increased upon request; refer to [AWS service quotas](#) for more information.

If you enable job queueing but don't exceed the quota for concurrent requests, all requests are processed concurrently.

Enabling job queueing

You can enable job queueing using the **AWS Management Console**, **AWS SDK**, or **AWS CLI**; see the following for examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This opens the **Specify job details** page.
3. In the **Job Settings** box, there is an **Additional settings** panel. If you expand this panel, you can select the **Add to job queue** box to enable job queueing.

Specify job details Info

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Language settings

You can transcribe your audio file in a language that you specify or have Amazon Transcribe identify and transcribe it in the predominant language.

Specific language Info

If you know the language spoken in your source audio, choose this option to get the most accurate results. The options available for additional processing vary between languages.

Automatic language identification Info

If you don't know the language spoken in your audio files, choose this option. You have access to fewer options for additional processing than if you choose **Specific language**.

Language

Choose the language of the input audio.

Model type Info

Choose the type of model to use for the transcription job.

General model

To use a model that is not specialized for a particular use case, choose this option. Configuration options vary between languages.

Custom language model

To use a model that you trained for your specific use case, choose this option. This model has fewer configuration options than the general model.

▼ **Additional settings**

Job queue - optional Info

Enables you to submit jobs beyond the limit for concurrent jobs (100). You must specify access permissions to the resources that job queuing uses.

Add to job queue

4. Fill in any other fields you wish to include on the **Specify job details** page, then select **Next**. This takes you to the **Configure job - optional page**.
5. Select **Create job** to run your transcription job.

AWS CLI

This example uses the `start-transcription-job` command and `job-execution-settings` parameter with the `AllowDeferredExecution` sub-parameter. Note that when you include `AllowDeferredExecution` in your request, you must also include `DataAccessRoleArn`.

For more information, see [StartTranscriptionJob](#) and [JobExecutionSettings](#).

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--job-execution-settings AllowDeferredExecution
```

```
--language-code en-US \  
--job-execution-settings  
  AllowDeferredExecution=true,DataAccessRoleArn=arn:aws:iam::111122223333:role/ExampleRole
```

Here's another example using the `start-transcription-job` command, and a request body that adds subtitles to that job.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://my-first-queueing-request.json
```

The file `my-first-queueing-request.json` contains the following request body.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "LanguageCode": "en-US",  
  "JobExecutionSettings": {  
    "AllowDeferredExecution": true,  
    "DataAccessRoleArn": "arn:aws:iam::111122223333:role/ExampleRole"  
  }  
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to enable job queueing using the `AllowDeferredExecution` argument for the `start_transcription_job` method. Note that when you include `AllowDeferredExecution` in your request, you must also include `DataAccessRoleArn`. For more information, see [StartTranscriptionJob](#) and [JobExecutionSettings](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe', 'us-west-2')  
job_name = "my-first-queueing-request"  
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
transcribe.start_transcription_job(  
    TranscriptionJobName = job_name,  
    Media = {  
        'MediaFileUri': job_uri  
    },  
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',  
    OutputKey = 'my-output-files/',  
    LanguageCode = 'en-US',  
    JobExecutionSettings = {  
        'AllowDeferredExecution': True,  
        'DataAccessRoleArn': 'arn:aws:iam::111122223333:role/ExampleRole'  
    }  
)  
  
while True:  
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)  
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:
```

```
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

You can view the progress of a queued job via the AWS Management Console or by submitting a [GetTranscriptionJob](#) request. When a job is queued, the `Status` is `QUEUED`. The status changes to `IN_PROGRESS` once your job starts processing, then changes to either `COMPLETED` or `FAILED` when processing is finished.

Tagging resources

A tag is a custom metadata label that you can add to a resource in order to make it easier to identify, organize, and find in a search. Tags are comprised of two individual parts: A tag key and a tag value. This is referred to as a key:value pair.

A tag key typically represents a larger category, while a tag value represents a subset of that category. For example you could have *tag key=Color* and *tag value=Blue*, which would produce the key:value pair `Color:Blue`. Note that you can set the value of a tag to an empty string, but you can't set the value of a tag to null. Omitting the tag value is the same as using an empty string.

Tip

AWS Billing and Cost Management can use tags to separate your bills into dynamic categories. For example, if you add tags to represent different departments within your company, such as `Department:Sales` or `Department:Legal`, AWS can provide you with your cost distribution per department.

In Amazon Transcribe, you can tag the following resources:

- Transcription jobs
- Medical transcription jobs
- Custom vocabularies
- Custom medical vocabularies
- Custom vocabulary filters
- Custom language models

Tag keys can be up to 128 characters in length and tag values can be up to 256 characters in length; both are case sensitive. Amazon Transcribe supports up to 50 tags per resource. For a given resource, each tag key must be unique with only one value. Note that your tags cannot begin with `aws:` because AWS reserves this prefix for system-generated tags. You cannot add, modify, or delete `aws:*` tags, and they don't count against your tags-per-resource limit.

API operations specific to resource tagging

[ListTagsForResource](#), [TagResource](#), [UntagResource](#)

To use the tagging APIs, you need to include an Amazon Resource Name (ARN) with your request. ARNs have the format `arn:partition:service:region:account-id:resource-type/resource-id`. For example, the ARN associated with a transcription job may look like: `arn:aws:transcribe:us-west-2:111122223333:transcription-job/my-transcription-job-name`.

To learn more about tagging, including best practices, see [Tagging AWS resources](#).

Tag-based access control

You can use tags to control access within your AWS accounts. For tag-based access control, you provide tag information in the condition element of an IAM policy. You can then use tags and their associated tag condition key to control access to:

- **Resources:** Control access to your Amazon Transcribe resources based on the tags you've assigned to those resources.
 - Use the `aws:ResourceTag/key-name` condition key to specify which tag key:value pair must be attached to the resource.

- **Requests:** Control which tags can be passed in a request.
 - Use the `aws:RequestTag/key-name` condition key to specify which tags can be added, modified, or removed from an IAM user or role.
- **Authorization processes:** Control tag-based access for any part of your authorization process.
 - Use the `aws:TagKeys/` condition key to control whether specific tag keys can be used on a resource, in a request, or by a principal. In this case, the key value doesn't matter.

For an example tag-based access control policy, see [Viewing transcription jobs based on tags \(p. 239\)](#).

For more detailed information on tag-based access control, see [Controlling access to AWS resources using tags](#).

Adding tags to your Amazon Transcribe resources

You can add tags before or after you run your Amazon Transcribe job. The existing **Create*** and **Start*** APIs allow you to add tags with your transcription job.

You can add, modify or delete tags using the **AWS Management Console**, **AWS CLI**, or **AWS SDK**; see the following for examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This opens the **Specify job details** page.
3. Scroll to the bottom of the **Specify job details** page to find the **Tags - optional** box and select **Add new tag**.

Tags - optional
A tag is a label that you can add to a resource as metadata to help you organize, search, or filter your data. Each tag consists of a key and an option value.

No tags associated with the resource.

You can add up to 50 more tags.

4. Enter information for the **Key** field and, optionally, the **Value** field.

Tags - optional
A tag is a label that you can add to a resource as metadata to help you organize, search, or filter your data. Each tag consists of a key and an option value.

Key:


Value - optional:

You can add up to 49 more tags.

5. Fill in any other fields you wish to include on the **Specify job details** page, then select **Next**. This takes you to the **Configure job - optional** page.

Select **Create job** to run your transcription job.

6. You can view the tags associated with a transcription job by navigating to the **Transcription jobs** page, selecting a transcription job, and scrolling to the bottom of that job's information page. If you wish to edit your tags, you can do so by selecting **Manage tags**.



AWS CLI

This example uses the `start-transcription-job` command and `Tags` parameter. For more information, see [StartTranscriptionJob](#) and [Tag](#).

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--language-code en-US \  
--tags Key=color,Value=blue Key=shape,Value=square
```

Here's another example using the `start-transcription-job` command, and a request body that adds tags to that job.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://filepath/my-first-tagging-job.json
```

The file `my-first-tagging-job.json` contains the following request body.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "LanguageCode": "en-US",  
  "Tags": [  
    {  
      "Key": "color",  
      "Value": "blue"  
    },  
    {  
      "Key": "shape",  
      "Value": "square"  
    }  
  ]  
}
```


AWS SDK for Python (Boto3)

The following example uses the AWS SDK for Python (Boto3) to add a tag by using the `Tags` argument for the `start_transcription_job` method. For more information, see [StartTranscriptionJob](#) and [Tag](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
transcribe.start_transcription_job(
    TranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'my-output-files/',
    LanguageCode = 'en-US',
    Tags = [
        {
            'Key': 'color',
            'Value': 'blue'
        }
    ]
)

while True:
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Improving transcription accuracy with custom vocabularies and custom language models

If your media contains domain-specific or non-standard terms, such as brand names, acronyms, technical words, and jargon, you may notice that Amazon Transcribe doesn't correctly capture these terms in your transcription output.

To correct transcription inaccuracies and customize your output for your specific use case, you can create [Custom vocabularies \(p. 101\)](#) and [Custom language models \(p. 115\)](#).

- [Custom vocabularies \(p. 101\)](#) are designed to tune and boost both the recognition and the formatting of specific words in all contexts. This involves supplying Amazon Transcribe with words and, optionally, pronunciation and display forms.

If Amazon Transcribe is not correctly rendering specific terms in your transcripts, you can create a custom vocabulary file that tells Amazon Transcribe how you want these terms displayed. This word-specific approach is most appropriate for correcting terms like brand names and acronyms.

- [Custom language models \(p. 115\)](#) are designed to capture the context associated with terms. This involves supplying Amazon Transcribe with a large volume of domain-specific text data.

If Amazon Transcribe is not correctly rendering technical terms or is using the incorrect homophone in your transcripts, you can create a custom language model that teaches Amazon Transcribe your domain-specific language. For example, a custom language model can learn when to use 'floe' (ice floe) versus 'flow' (linear flow).

This context-aware approach is most appropriate for transcribing large volumes of domain-specific speech. Custom language models can produce significant accuracy improvements over custom vocabularies alone. When using batch transcriptions, you can include both a custom language model and a custom vocabulary in your request.

Tip

To achieve the highest transcription accuracy, use custom vocabularies in conjunction with your custom language models.

For a video demo on how to create a custom vocabulary using the AWS Management Console, see [Using a custom vocabulary](#).

For a video demo on how to create and use custom language models, see [Using Custom Language Models \(CLM\) to supercharge transcription accuracy](#).

Dive deeper with the AWS Machine Learning Blog

Learn how Formula 1 is using custom vocabularies to create real-time captions for their live races: [Live transcriptions of F1 races using Amazon Transcribe](#).

[Boost transcription accuracy of class lectures with custom language models for Amazon Transcribe](#).

Custom vocabularies

Use custom vocabularies to improve transcription accuracy for one or more specific words. These are generally domain-specific terms, such as brand names and acronyms, proper nouns, and words that Amazon Transcribe isn't rendering correctly.

Important

You are responsible for the integrity of your own data when you use Amazon Transcribe. Do not enter confidential information, personal information (PII), or protected health information (PHI) into a custom vocabulary.

Considerations when creating a custom vocabulary:

- You can have up to 100 vocabularies per AWS account
- The size limit for each custom vocabulary is 50 Kb
- Each entry cannot exceed 256 characters
- Only the characters listed in your language's [character set \(p. 5\)](#) can be used
- Vocabularies can be in table or list format (table format is strongly recommended)
- How you create your custom vocabulary varies depending on the format you used; refer to [Using a custom vocabulary \(p. 110\)](#) for examples
 - Table format: You must first upload your custom vocabulary file into an Amazon S3 bucket, then include the Amazon S3 URI of your vocabulary file in your request
 - List format: If using the AWS Management Console, you must save your custom vocabulary file locally, then include the path to your custom vocabulary file in your request; if using the command line (AWS CLI or AWS SDK), you must include your list as comma-separated words within your request

Tip

You can test your custom vocabulary using the AWS Management Console. Once your custom vocabulary is ready to use, log into the AWS Management Console, select **Real-time transcription**, scroll to **Customizations**, toggle on **Custom vocabulary**, and select your custom vocabulary from the drop-down list. Then select **start streaming**. Speak some of the words in your custom vocabulary into your microphone to see if they render correctly.

Custom vocabulary tables versus lists

Tables are strongly preferred.

Tables give you more options for—and more control over—the input and output of words within your custom vocabulary. With tables, you must specify multiple categories (Phrase, IPA, SoundsLike, and DisplayAs), allowing you to fine-tune your output.

Lists don't have additional options, so you can only type in entries as you want them to appear in your transcript, replacing all spaces with hyphens. Custom vocabulary lists are not supported with the [CreateMedicalVocabulary](#) operation.

The AWS Management Console, AWS CLI, and AWS SDKs all use custom vocabulary tables in the same way; lists are used differently for each method and thus may require additional formatting for successful use between methods.

For more information, see [Creating a custom vocabulary using a table \(p. 102\)](#) and [Creating a custom vocabulary using a list \(p. 106\)](#).

To dive a little deeper and learn how to use Amazon Augmented AI with custom vocabularies, see [Start building a human review along with Amazon Transcribe](#)

API operations specific to custom vocabularies

[CreateVocabulary](#), [DeleteVocabulary](#), [GetVocabulary](#), [ListVocabularies](#), [UpdateVocabulary](#)

Creating a custom vocabulary using a table

Using a table format is the most robust way to create your custom vocabulary. Vocabulary tables must consist of four columns, which can be included in any order:

Phrase	SoundsLike	IPA	DisplayAs
Required. Every row in your table must contain an entry in this column.	Optional. Rows in this column can be left empty.	Optional. Rows in this column can be left empty.	Optional. Rows in this column can be left empty.
Do not use spaces in this column.	Do not use spaces in this column.	You must add a single space between every IPA character (single-byte) or valid IPA character pair (double-byte).	You can use spaces in this column.
If your entry contains multiple words, separate each word with a hyphen (-). For example, Andorra-la-Vella or Los-Angeles .	Include your entry as hyphen-separated syllables that mimic how the word sounds. It's better to use small, common words over phonetic representations. For example, for 'Los Angeles', ' loss-ann-ge1-es ' is preferable to ' lahs-ahn-jul-ees '.	This column is intended for phonetic spellings using only characters in the International Phonetic Alphabet (IPA) . For example, 'Los Angeles' is l # s æ n # # l # s and 'CLI' is s # # l a# .	Defines the how you want your entry to look in your transcription output. For example, Andorra-la-Vella in the <code>Phrase</code> column is Andorra la Vella in the <code>DisplayAs</code> column.
For acronyms, any pronounced letters must be separated by a period. If your acronym is plural, you must use a hyphen between the acronym and the 's'. For example, 'CLI' is C.L.I. and 'ABCs' is A.B.C.-s .	You cannot have an entry for both <code>SoundsLike</code> and <code>IPA</code> in the same row.	You cannot have an entry for both <code>IPA</code> and <code>SoundsLike</code> in the same row.	If a row in this column is empty, Amazon Transcribe uses the contents of the <code>Phrase</code> column to determine output.
If your phrase consists of both a word and an acronym, these two components must be separated by a hyphen. For example, 'DynamoDB' is Dynamo-D.B..			

Things to note when creating your table:

- Your table must contain all four columns (`Phrase`, `SoundsLike`, `IPA`, and `DisplayAs`), but the `Phrase` column is the only one that must contain an entry on each row. All other columns can be left empty.
- In a given row, you cannot have entries for both `IPA` and `SoundsLike` fields. You must choose one or the other, or leave both blank.
- You can only use characters that are supported for your language. Refer to your language's [character set \(p. 5\)](#) for details.
- Spaces are only allowed within the `IPA` and `DisplayAs` columns. Do not use spaces to separate columns; columns must be TAB-separated.

- You must save your table as a plain text (*.txt) file in LF format. If you use any other format, such as CRLF, your custom vocabulary is not accepted by Amazon Transcribe.
- You must upload your custom vocabulary file into an Amazon S3 bucket and process it using [CreateVocabulary](#) before you can include it in a transcription request. Refer to [Creating custom vocabulary tables \(p. 103\)](#) for instructions.

Note

Enter acronyms, or other words whose letters should be pronounced individually, as single letters separated by periods (A.B.C.). To enter the plural form of an acronym, such as 'ABCs', separate the 's' from the acronym with a hyphen (A.B.C.-s). You can use upper or lower case letters to define an acronym. Acronyms are not supported in all languages; refer to [Supported languages and language-specific features \(p. 3\)](#).

Here is a sample custom vocabulary table (where [TAB] represents a tab character):

```
Phrase[TAB]SoundsLike[TAB]IPA[TAB]DisplayAs
Los-Angeles[TAB][TAB]l # s æ n # # l # s[TAB]Los Angeles
Eva-Maria[TAB]ay-va-ma-ree-ah[TAB][TAB]
A.B.C.-s[TAB]ay-bee-sees[TAB][TAB]ABCs
Amazon-dot-com[TAB][TAB][TAB]Amazon.com
C.L.I.[TAB][TAB]s # # l a#[TAB]CLI
Andorra-la-Vella[TAB]ann-do-rah-la-bay-ah[TAB][TAB]Andorra la Vella
Dynamo-D.B.[TAB][TAB][TAB]DynamoDB
```

Here is the same table with aligned columns for visual clarity. **Do not** add spaces between columns in your custom vocabulary table; your table should look misaligned like the preceding example.

```
Phrase           [TAB]SoundsLike           [TAB]IPA           [TAB]DisplayAs
Los-Angeles     [TAB]                       [TAB]l # s æ n # # l # s[TAB]Los Angeles
Eva-Maria       [TAB]ay-va-ma-ree-ah      [TAB]               [TAB]
A.B.C.-s        [TAB]ay-bee-sees          [TAB]               [TAB]ABCs
amazon-dot-com  [TAB]                       [TAB]               [TAB]amazon.com
C.L.I.          [TAB]                       [TAB]s # # l a#      [TAB]CLI
Andorra-la-Vella[TAB]ann-do-rah-la-bay-ah[TAB] [TAB]Andorra la Vella
Dynamo-D.B.    [TAB]                       [TAB]               [TAB]DynamoDB
```


Creating custom vocabulary tables

To process a custom vocabulary table for use with Amazon Transcribe, see the following examples:

AWS Management Console

Before continuing, save your custom vocabulary as a text (*.txt) file, then upload it into an Amazon S3 bucket.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Custom vocabulary**. This opens the **Custom vocabulary** page where you can view existing vocabularies or create a new one.
3. Select **Create vocabulary**.



This takes you to the **Create vocabulary** page. Enter a name for your new custom vocabulary.

Select the **S3 location** option under **Vocabulary input source**. Then, either manually enter the Amazon S3 path or select **Browse S3** to locate your custom vocabulary.

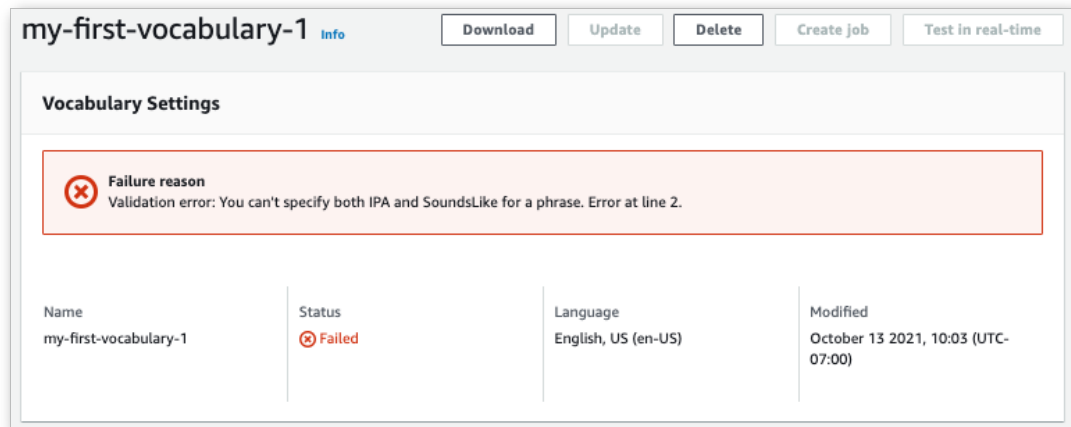
- Optionally, add tags to your custom vocabulary. Once you have all fields completed, select **Create vocabulary** at the bottom of the page. This takes you back to the **Custom vocabulary** page where you can view the status of your custom vocabulary. When the status changes from 'Pending' to 'Ready' your custom vocabulary can be used with a transcription.

Name	Language	Status
my-first-vocabulary-1	English, US (en-US)	Pending

- If the status changes to 'Failed', click on the name of your custom vocabulary to go to its information page.

Name	Language	Status
my-first-vocabulary-2	English, US (en-US)	Ready
my-first-vocabulary-1	English, US (en-US)	Failed

There is a **Failure reason** banner at the top of this page that provides information on why your custom vocabulary failed. Correct the error in your text file and try again.



AWS CLI

This example uses the `create-vocabulary` command with a table-formatted vocabulary file. For more information, see [CreateVocabulary](#).

To use an existing custom vocabulary in a transcription job, set the `VocabularyName` in the `Settings` field when you call the `StartTranscriptionJob` operation or, from the AWS Management Console, choose the custom vocabulary from the drop-down list.

```
aws transcribe create-vocabulary \  
--vocabulary-name my-first-vocabulary \  
--vocabulary-file-uri s3://DOC-EXAMPLE-BUCKET/my-vocabularies/my-vocabulary-file.txt \  
--language-code en-US
```

Here's another example using the `create-vocabulary` command, and a request body that creates your vocabulary.

```
aws transcribe create-vocabulary \  
--cli-input-json file://filepath/my-first-vocab-table.json
```

The file `my-first-vocab-table.json` contains the following request body.

```
{  
  "VocabularyName": "my-first-vocabulary",  
  "VocabularyFileUri": "s3://DOC-EXAMPLE-BUCKET/my-vocabularies/my-vocabulary-table.txt",  
  "LanguageCode": "en-US"  
}
```

Once `VocabularyState` changes from `PENDING` to `READY`, your custom vocabulary is ready to use with a transcription. To view the current status of your custom vocabulary, run:

```
aws transcribe get-vocabulary \  
--vocabulary-name my-first-vocabulary
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to create a custom vocabulary from a table using the `create_vocabulary` method. For more information, see [CreateVocabulary](#).

To use an existing custom vocabulary in a transcription job, set the `VocabularyName` in the [Settings](#) field when you call the [StartTranscriptionJob](#) operation or, from the AWS Management Console, choose the custom vocabulary from the drop-down list.

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
vocab_name = "my-first-vocabulary"
response = transcribe.create_vocabulary(
    LanguageCode = 'en-US',
    VocabularyName = vocab_name,
    VocabularyFileUri = 's3://DOC-EXAMPLE-BUCKET/my-vocabularies/my-vocabulary-table.txt'
)

while True:
    status = transcribe.get_vocabulary(VocabularyName = vocab_name)
    if status['VocabularyState'] in ['READY', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Note

If you create a new Amazon S3 bucket for your custom vocabulary files, make sure the IAM role making the [CreateVocabulary](#) request has permissions to access this bucket. If the role doesn't have the correct permissions, your request fails. You can optionally specify an IAM role within your request by including the `DataAccessRoleArn` parameter. For more information on IAM roles and policies in Amazon Transcribe, see [Amazon Transcribe identity-based policy examples \(p. 235\)](#).

Creating a custom vocabulary using a list

You can create custom vocabularies from lists using the AWS Management Console, AWS CLI, or AWS SDKs.

- **AWS Management Console:** You must create and upload a text file containing your custom vocabulary. You can use either line-separated or comma-separated entries. Note that your list must be saved as a text (*.txt) file in LF format. If you use any other format, such as CRLF, your custom vocabulary is not accepted by Amazon Transcribe.
- **AWS CLI and AWS SDKs:** You must include your custom vocabulary as comma-separated entries within your API call using the [Phrases](#) flag.

If an entry contains multiple words, you must hyphenate each word. For example, you include 'Los Angeles' as **Los-Angeles** and 'Andorra la Vella' as **Andorra-la-Vella**.

Here are examples of the two valid list formats. Refer to [Creating custom vocabulary lists \(p. 107\)](#) for method-specific examples.

- Comma-separated entries:

```
Los-Angeles,CLI,Eva-Maria,ABCs,Andorra-la-Vella
```

- Line-separated entries:


```
Los-Angeles  
CLI  
Eva-Maria  
ABCs  
Andorra-la-Vella
```

Important

You can only use characters that are supported for your language. Refer to your language's [character set \(p. 5\)](#) for details.

Custom vocabulary lists are not supported with the [CreateMedicalVocabulary](#) operation. If creating a custom medical vocabulary, you must use a table format; refer to [Creating a custom vocabulary using a table \(p. 102\)](#) for instructions.


Creating custom vocabulary lists

To process a custom vocabulary list for use with Amazon Transcribe, see the following examples:

AWS Management Console

Before continuing, save your custom vocabulary as a text (*.txt) file.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Custom vocabulary**. This opens the **Custom vocabulary** page where you can view existing vocabularies or create a new one.
3. Select **Create vocabulary**.



This takes you to the **Create vocabulary** page. Enter a name for your new custom vocabulary.

Select the **File upload** option under **Vocabulary input source**. Then, select **Choose file** to upload your custom vocabulary.

Create vocabulary [Info](#)

Vocabulary settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, and 0-9.

Language

Vocabulary input source [Info](#)

File upload
List format only

S3 location
Table format only

Vocabulary file

File format: txt, maximum size 50 KB.

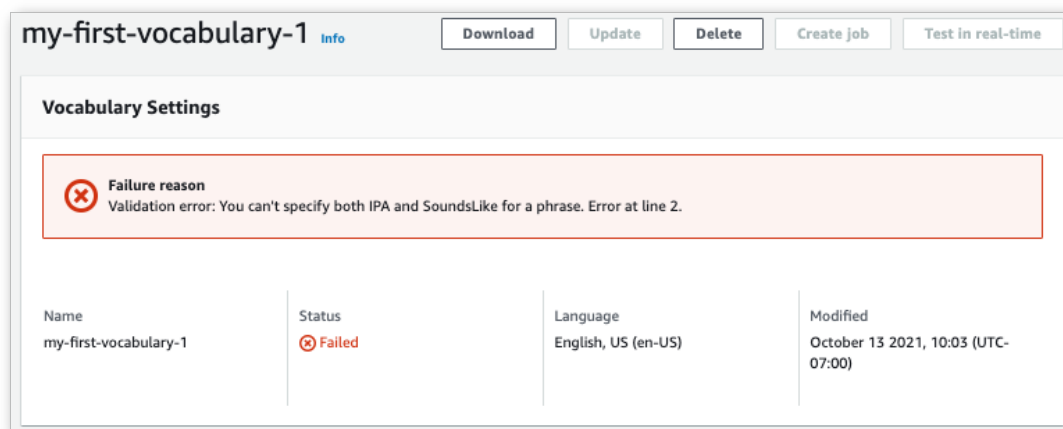
- Optionally, add tags to your custom vocabulary. Once you have all fields completed, select **Create vocabulary** at the bottom of the page. This takes you back to the **Custom vocabulary** page where you can view the status of your custom vocabulary. When the status changes from 'Pending' to 'Ready' your custom vocabulary can be used with a transcription.

Name	Language	Status
<input type="radio"/> my-first-vocabulary-1	English, US (en-US)	<input type="radio"/> Pending

- If the status changes to 'Failed', click on the name of your custom vocabulary to go to its information page.

Name	Language	Status
<input type="radio"/> my-first-vocabulary-2	English, US (en-US)	<input checked="" type="radio"/> Ready
<input type="radio"/> my-first-vocabulary-1	English, US (en-US)	<input checked="" type="radio"/> Failed

There is a **Failure reason** banner at the top of this page that provides information on why your custom vocabulary failed. Correct the error in your text file and try again.



AWS CLI

This example uses the `create-vocabulary` command with a list-formatted custom vocabulary file. For more information, see [CreateVocabulary](#).

```
aws transcribe create-vocabulary \  
--vocabulary-name my-first-vocabulary \  
--language-code en-US \  
--phrases {CLI,Eva-Maria,ABCs}
```

Here's another example using the `create-vocabulary` command, and a request body that creates your vocabulary.

```
aws transcribe create-vocabulary \  
--cli-input-json file://filepath/my-first-vocab-list.json
```

The file `my-first-vocab-list.json` contains the following request body.

```
{  
  "VocabularyName": "my-first-vocabulary",  
  "LanguageCode": "en-US",  
  "Phrases": [  
    "CLI", "Eva-Maria", "ABCs"  
  ]  
}
```

Once `VocabularyState` changes from `PENDING` to `READY`, your custom vocabulary is ready to use with a transcription. To view the current status of your custom vocabulary, run:

```
aws transcribe get-vocabulary \  
--vocabulary-name my-first-vocabulary
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to create a custom vocabulary from a list using the `create_vocabulary` method. For more information, see [CreateVocabulary](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
vocab_name = "my-first-vocabulary"
response = transcribe.create_vocabulary(
    LanguageCode = 'en-US',
    VocabularyName = vocab_name,
    Phrases = [
        'CLI', 'Eva-Maria', 'ABCs'
    ]
)

while True:
    status = transcribe.get_vocabulary(VocabularyName = vocab_name)
    if status['VocabularyState'] in ['READY', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Note

If you create a new Amazon S3 bucket for your custom vocabulary files, make sure the IAM role making the [CreateVocabulary](#) request has permissions to access this bucket. If the role doesn't have the correct permissions, your request fails. You can optionally specify an IAM role within your request by including the `DataAccessRoleArn` parameter. For more information on IAM roles and policies in Amazon Transcribe, see [Amazon Transcribe identity-based policy examples](#) (p. 235).

Using a custom vocabulary

Once your custom vocabulary is created, you can include it in your transcription requests; refer to the following sections for examples.

The language of the custom vocabulary you're including in your request must match the language code you specify for your media. If the languages don't match, your custom vocabulary is not applied to your transcription and there are no warnings or errors.

Using a custom vocabulary in a batch transcription

To use a custom vocabulary with a batch transcription, see the following for examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This opens the **Specify job details** page.

Specify job details [Info](#)

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Model type [Info](#)
Choose the type of model to use for the transcription job.

General model
To use a model that is not specialized for a particular use case, choose this option. Configuration options vary between languages.

Custom language model
To use a model that you trained for your specific use case, choose this option. This model has fewer configuration options than the general model.

Language settings
You can transcribe your audio file in a language that you specify or have Amazon Transcribe Identify and transcribe it in the predominant language.

Specific language [Info](#)
If you know the language spoken in your source audio, choose this option to get the most accurate results. The options available for additional processing vary between languages.

Automatic language identification [Info](#)
If you don't know the language spoken in your audio files, choose this option. You have access to fewer options for additional processing than if you choose **Specific language**.

Language
Choose the language of the input audio.

► **Additional settings**

- Name your job and specify your input media. Optionally include any other fields, then choose **Next**.
3. At the bottom of the **Configure job** page, in the **Customization** panel, toggle on **Custom vocabulary**.

Configure job - optional [Info](#)

Audio settings

Audio identification [Info](#)
Choose to split multi-channel audio into separate channels for transcription, or identify speakers in the input audio.

Alternative results [Info](#)
Enable to view more transcription results

Content removal

Content removal conceals information in the resulting transcript from your source audio file. Amazon Transcribe changes items in the transcript and does not modify the source audio.

PII redaction [Info](#)
Label the type of PII and also mask the content with the PII entity type in the transcription output. For example, (123) 456-7890 will be masked as [PHONE].

Vocabulary filtering [Info](#)
Vocabulary filtering can remove, mask or tag specified words in the final transcript.

Customization

Custom vocabulary [Info](#)
A custom vocabulary improves the accuracy of recognizing words and phrases specific to your use case.

Vocabulary selection
The vocabularies shown here are based on your language settings. You can choose up to one vocabulary per language. You can also [create a new vocabulary](#). [↗](#)

Choose a vocabulary ▼

Cancel Previous **Create job**

4. Select your custom vocabulary from the drop-down menu.

Select **Create job** to run your transcription job.

AWS CLI

This example uses the `start-transcription-job` command and `Settings` parameter with the `VocabularyName` sub-parameter. For more information, see [StartTranscriptionJob](#) and [Settings](#).

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--language-code en-US \  

```

```
--settings VocabularyName=my-first-vocabulary
```

Here's another example using the `start-transcription-job` command, and a request body that includes your custom vocabulary with that job.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://my-first-vocabulary-job.json
```

The file `my-first-vocabulary-job.json` contains the following request body.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "LanguageCode": "en-US",  
  "Settings": {  
    "VocabularyName": "my-first-vocabulary"  
  }  
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to include a custom vocabulary using the `Settings` argument for the `start_transcription_job` method. For more information, see [StartTranscriptionJob](#) and [Settings](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

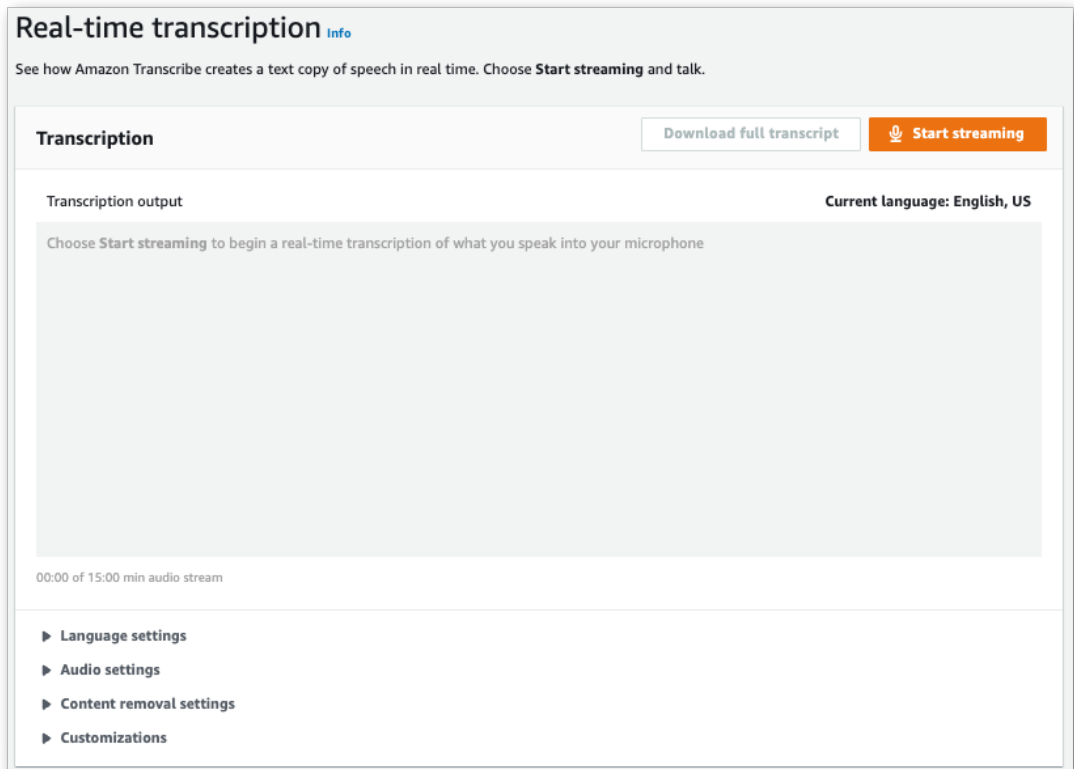
```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe', 'us-west-2')  
job_name = "my-first-transcription-job"  
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
transcribe.start_transcription_job(  
    TranscriptionJobName = job_name,  
    Media = {  
        'MediaFileUri': job_uri  
    },  
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',  
    OutputKey = 'my-output-files/',  
    LanguageCode = 'en-US',  
    Settings = {  
        'VocabularyName': 'my-first-vocabulary'  
    }  
)  
  
while True:  
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)  
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:  
        break  
    print("Not ready yet...")  
    time.sleep(5)  
print(status)
```

Using a custom vocabulary in a streaming transcription


To use a custom vocabulary with a streaming transcription, see the following for examples:

AWS Management Console

1. Sign into the [AWS Management Console](#).
2. In the navigation pane, choose **Real-time transcription**. Scroll down to **Customizations** and expand this field if it is minimized.



3. Toggle on **Custom vocabulary** and select a custom vocabulary from the drop-down menu.



Include any other settings you wish to apply to your stream.

4. You're now ready to transcribe your stream. Select **Start streaming** and begin speaking. To end your dictation, select **Stop streaming**.

HTTP/2 stream

This example creates an HTTP/2 request that includes your custom vocabulary. For more information on using HTTP/2 streaming with Amazon Transcribe, see [Setting up an HTTP/2 stream \(p. 82\)](#). For more detail on parameters and headers specific to Amazon Transcribe, see [StartStreamTranscription](#).


```
POST /stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
X-Amz-Target: com.amazonaws.transcribe.Transcribe.StartStreamTranscription
Content-Type: application/vnd.amazon.eventstream
X-Amz-Content-Sha256: string
X-Amz-Date: 20220208T235959Z
Authorization: AWS4-HMAC-SHA256 Credential=access-key/20220208/us-west-2/transcribe/
aws4_request, SignedHeaders=content-type;host;x-amz-content-sha256;x-amz-date;x-amz-
target;x-amz-security-token, Signature=string
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
x-amzn-transcribe-vocabulary-name: my-first-vocabulary
transfer-encoding: chunked
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

WebSocket stream

This example creates a pre-signed URL that applies your custom vocabulary to a WebSocket stream. Line breaks have been added for readability. For more information on using WebSocket streams with Amazon Transcribe, see [Setting up a WebSocket stream \(p. 86\)](#). For more detail on parameters, see [StartStreamTranscription](#).

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/stream-transcription-websocket?
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=string
&X-Amz-SignedHeaders=content-type%3Bhost%3Bx-amz-date
&language-code=en-US
&media-encoding=flac
&sample-rate=16000
&vocabulary-name=my-first-vocabulary
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

Custom language models

Custom language models are designed to improve transcription accuracy for domain-specific speech. This includes any content outside of what you would hear in normal, everyday conversations. For example, if you're transcribing the proceedings from a scientific conference, a standard transcription is unlikely to recognize many of the scientific terms used by presenters. In this case, you can train a custom language model to recognize the specialized terms used in your discipline.

Unlike custom vocabularies, which increase recognition of a word by providing hints (such as pronunciations), custom language models learn the context associated with a given word. This includes how and when a word is used, and the relationship a word has to other words. For example, if you train your model using climate science research papers, your model may learn that 'ice floe' is a more likely word pair than 'ice flow'.

To view the supported languages for custom language models, refer to [Supported languages and language-specific features \(p. 3\)](#).

API operations specific to custom language models

[CreateLanguageModel](#), [DeleteLanguageModel](#), [DescribeLanguageModel](#), [ListLanguageModels](#)

Data sources

You can use any type of text data you want to train your model. However, the closer your text content is to your audio content, the more accurate your model. Therefore, it's important to choose text data that use the same terms in the same context as your audio.

The best data for training a model are accurate transcripts. This is considered in-domain data. In-domain text data have the exact same terms, usage, and context as the audio you want to transcribe.

If you don't have accurate transcripts, use journal articles, technical reports, whitepapers, conference proceedings, instruction manuals, news articles, website content, and any other text that contains the desired terms used in a similar context to that of your audio. This is considered domain-related data.

Creating a robust custom language model may require a significant amount of text data, which must contain the terms spoken in your audio. You can supply Amazon Transcribe with up to 2 GB of text data to train your model—this is referred to as **training data**. Optionally, when you have no (or few) in-domain transcripts, you can provide Amazon Transcribe with up to 200 MB of text data to tune your model—this is referred to as **tuning data**.

Training versus tuning data

The purpose of training data is to teach Amazon Transcribe to recognize new terms and learn the context in which these terms are used. In order to create a robust model, Amazon Transcribe may require a large volume of relevant text data. Providing as much training data as possible, up to the 2 GB limit, is strongly recommended.

The purpose of tuning data is to help refine and optimize the contextual relationships learned from your training data. Tuning data are not required to create a custom language model.

It's up to you to decide how best to select training and, optionally, tuning data. Each case is unique and depends on the type and amount of data you have. Tuning data are recommended when you lack in-domain training data.

If you choose to include both data types, do **not** overlap your training and tuning data; training and tuning data should be unique. Overlapping data can bias and skew your custom language model, impacting its accuracy.

As general guidance, we recommend using accurate, in-domain text as training data whenever possible. Here are some general scenarios, listed in order of preference:

- If you have more than 10,000 words of accurate, in-domain transcript text, use it as training data. In this case, there is no need to include tuning data. This is the ideal scenario for training a custom language model.
- If you have accurate, in-domain transcript text that contains less than 10,000 words and are not getting the desired results, consider augmenting your training data with domain-related written texts, such as technical reports. In this case, reserve a small portion (10-25%) of your in-domain transcript data to use as tuning data.
- If you have no in-domain transcript text, upload all your domain-related text as training data. In this case, transcript-style text is preferable to written text. This is the least effective scenario for training a custom language model.

When you're ready to create your model, see [Creating a custom language model \(p. 117\)](#).

Creating a custom language model

Before you can create your custom language model, you must:

- Prepare your data. Data must be saved in plain text format and can't contain any special characters.
- Upload your data into an Amazon S3 bucket. Creating separate folders for training and tuning data is recommended.
- Make sure Amazon Transcribe has access to your Amazon S3 bucket. You must specify an IAM role that has access permissions to use your data.

Preparing your data

You can compile all your data in one file or save it as multiple files. Note that if you choose to include tuning data, it must be saved in a separate file from your training data.

It doesn't matter how many text files you use for your training or tuning data. Uploading one file with 100,000 words produces the same result as uploading 10 files with 10,000 words. Prepare your text data in a way that's most convenient for you.

Make sure all your data files meet the following criteria:

- They are all in the same language as the model you want to create. For example, if you want to create a custom language model that transcribes audio in US English (en-US), all your text data must be in US English.
- They are in plain text format with UTF-8 encoding.
- They don't contain any special characters or formatting, such as HTML tags.
- They amount to a maximum combined total of 2 GB in size for training data and 200 MB for tuning data.

If any of these criteria are not met, your model fails.

Uploading your data

Before uploading your data, create a new folder for your training data. If using tuning data, create another separate folder.

The URIs for your buckets might look like:

- `s3://DOC-EXAMPLE-BUCKET/my-model-training-data/`
- `s3://DOC-EXAMPLE-BUCKET/my-model-tuning-data/`

Upload your training and tuning data into the appropriate buckets.

You can add more data to these buckets at a later date. However, if you do, you need to recreate your model with the new data. Existing models can't be updated with new data.

Allowing access to your data

To create a custom language model, you must specify an IAM role that has permissions to access your Amazon S3 bucket. If you don't already have a role with access to the Amazon S3 bucket where you placed your training data, you must create one. After you create a role, you can attach a policy to grant that role permissions.

For example policies, see [Amazon Transcribe identity-based policy examples \(p. 235\)](#).

To learn how to create a new IAM role, see [Creating your first IAM admin user and user group](#) or [Creating your first IAM delegated user and user group](#).

To learn more about policies, see [Policies and permissions in IAM](#) and [Creating IAM policies](#).

Creating your custom language model


When creating your custom language model, you must choose a base model. There are two base model options:

- **NarrowBand**: Use this option for audio with a sample rate of less than 16,000 Hz. This model type is typically used for telephone conversations recorded at 8,000 Hz.
- **WideBand**: Use this option for audio with a sample rate greater than or equal to 16,000 Hz.

You can create custom language models using the AWS Management Console, AWS CLI, or AWS SDKs.; see the following examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Custom language model**. This opens the **Custom language models** page where you can view existing custom language models or train a new custom language model.
3. To train a new model, select **Train model**.



This takes you to the **Train model** page. Add a name, specify the language, and choose the base model you want for your model. Then, add the path to your training and, optionally, your tuning data. You must include an IAM role that has permissions to access your data.

Train model [Info](#)

Model settings

Name

The name can be up to 200 characters long. Valid characters:A-Z,a-z,0-9,and _ - (hyphen).

Language
Choose the language of your model.

Base model [Info](#)
Choose the base model that you want to use to create your custom language model. Choose the model based on the sample rate of your source audio.

Narrow band
For audio that has a sample rate less than 16 KHz. Typically, this is 8 KHz audio from telephone conversations.

Wide band
For audio that has a sample rate of 16 KHz or greater. Typically, this is 16 KHz audio from media sources.

Training data [Info](#)

Training data location on S3
Type or paste the S3 prefix for the text files that you want to use as training data, or browse to find the files that have matching S3 prefixes.

The file format must be plain text in the language that you have selected for the model. The maximum file size is 2 GB.

Tuning data - optional [Info](#)

Tuning data location on S3
Type or paste the S3 prefix for the text files that you want to use as tuning data, or browse to find the files that have matching S3 prefixes.

The file format must be plain text in the language that you have selected for the model. The maximum file size is 200 MB.

Access permissions

IAM role [Info](#)

Use an existing IAM role

Create an IAM role
By choosing Train model you are authorizing creation of this role.

Role name
A role that grants access to the S3 input locations.

- Once you have all fields completed, select **Train model** at the bottom of the page.

AWS CLI

This example uses the `create-language-model` command. For more information, see [CreateLanguageModel](#) and [LanguageModel](#).

```
aws transcribe create-language-model \  
--base-model-name NarrowBand \  
--model-name my-first-language-model \  
--input-data-config S3Uri=s3://DOC-EXAMPLE-BUCKET/my-clm-training-  
data/,TuningDataS3Uri=s3://DOC-EXAMPLE-BUCKET/my-clm-tuning-  
data/,DataAccessRoleArn=arn:aws:iam::11112223333:role/ExampleRole \  

```

```
--language-code en-US
```

Here's another example using the [create-language-model](#) command, and a request body that creates your CLM.

```
aws transcribe create-language-model \  
--cli-input-json file://filepath/my-first-language-model.json
```

The file *my-first-language-model.json* contains the following request body.

```
{  
  "BaseModelName": "NarrowBand",  
  "ModelName": "my-first-language-model",  
  "InputDataConfig": {  
    "S3Uri": "s3://DOC-EXAMPLE-BUCKET/my-clm-training-data/",  
    "TuningDataS3Uri": "s3://DOC-EXAMPLE-BUCKET/my-clm-tuning-data/",  
    "DataAccessRoleArn": "arn:aws:iam::111122223333:role/ExampleRole"  
  },  
  "LanguageCode": "en-US"  
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to create a CLM using the [create_language_model](#) method. For more information, see [CreateLanguageModel](#) and [LanguageModel](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe', 'us-west-2')  
model_name = 'my-first-language-model',  
transcribe.create_language_model(  
    LanguageCode = 'en-US',  
    BaseModelName = 'NarrowBand',  
    ModelName = model_name,  
    InputDataConfig = {  
        'S3Uri': 's3://DOC-EXAMPLE-BUCKET/my-clm-training-data/',  
        'TuningDataS3Uri': 's3://DOC-EXAMPLE-BUCKET/my-clm-tuning-data/',  
        'DataAccessRoleArn': 'arn:aws:iam::111122223333:role/ExampleRole'  
    }  
)  
  
while True:  
    status = transcribe.get_language_model(ModelName = model_name)  
    if status['LanguageModel']['ModelStatus'] in ['COMPLETED', 'FAILED']:  
        break  
    print("Not ready yet...")  
    time.sleep(5)  
    print(status)
```

Updating your custom language model

Amazon Transcribe continually updates the base models available for custom language models.

To see if your custom language model is using the latest base model, run a [DescribeLanguageModel](#) request using the AWS CLI or an AWS SDK, then find the `UpgradeAvailability` field in your response.

If `UpgradeAvailability` is `true`, your model is not running the latest version of the base model. To use the latest base model in a custom language model, you must create a new custom language model. Custom language models cannot be upgraded.

Using a custom language model

Once you've created your custom language model, you can include it in your transcription requests; refer to the following sections for examples.

The language of the model you're including in your request must match the language code you specify for your media. If the languages don't match, your custom language model is not applied to your transcription and there are no warnings or errors.

Using a custom language model in a batch transcription

To use a custom language model with a batch transcription, see the following for examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This opens the **Specify job details** page.
3. In the **Job settings** panel under **Model type**, select the **Custom language model** box.

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Model type [Info](#)
Choose the type of model to use for the transcription job.

General model
To use a model that is not specialized for a particular use case, choose this option. Configuration options vary between languages.

Custom language model
To use a model that you trained for your specific use case, choose this option. This model has fewer configuration options than the general model.

Language
Choose the language of the input audio.

Custom model selection
Choose an existing model or [create a new one](#). [↗](#)

▶ Additional settings

You must also select an input language from the drop-down menu.

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Model type [Info](#)
Choose the type of model to use for the transcription job.

General model
To use a model that is not specialized for a particular use case, choose this option. Configuration options vary between languages.

Custom language model
To use a model that you trained for your specific use case, choose this option. This model has fewer configuration options than the general model.

Language
Choose the language of the input audio.

- English, US (en-US)
- English, AU (en-AU)
- English, UK (en-GB)
- Hindi, IN (hi-IN)
- Spanish, US (es-US)

4. Under **Custom model selection**, either select an existing custom language model from the drop-down menu or select **Create a new one**.

Add the Amazon S3 location of your input file in the **Input data** panel.

5. Select **Next** for additional configuration options.

Select **Create job** to run your transcription job.

AWS CLI

This example uses the `start-transcription-job` command and `ModelSettings` parameter with the `VocabularyName` sub-parameter. For more information, see [StartTranscriptionJob](#) and [ModelSettings](#).

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--language-code en-US \  
--model-settings LanguageModelName=my-first-language-model
```

Here's another example using the `start-transcription-job` command, and a request body that includes your custom vocabulary with that job.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://my-first-model-job.json
```


The file *my-first-model-job.json* contains the following request body.

```
{
  "TranscriptionJobName": "my-first-transcription-job",
  "Media": {
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
  },
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",
  "OutputKey": "my-output-files/",
  "LanguageCode": "en-US",
  "ModelSettings": {
    "LanguageModelName": "my-first-language-model"
  }
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to include a custom language model using the `ModelSettings` argument for the `start_transcription_job` method. For more information, see [StartTranscriptionJob](#) and [ModelSettings](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
transcribe.start_transcription_job(
    TranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'my-output-files/',
    LanguageCode = 'en-US',
    ModelSettings = {
        'LanguageModelName': 'my-first-language-model'
    }
)

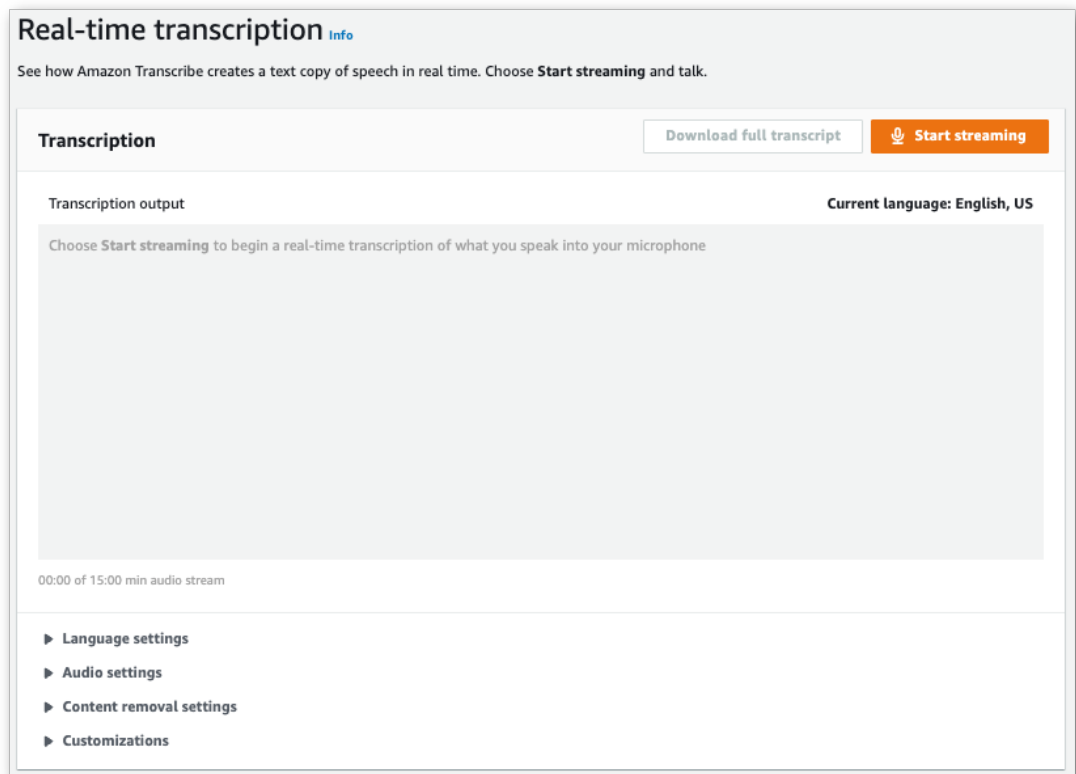
while True:
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Using a custom language model in a streaming transcription

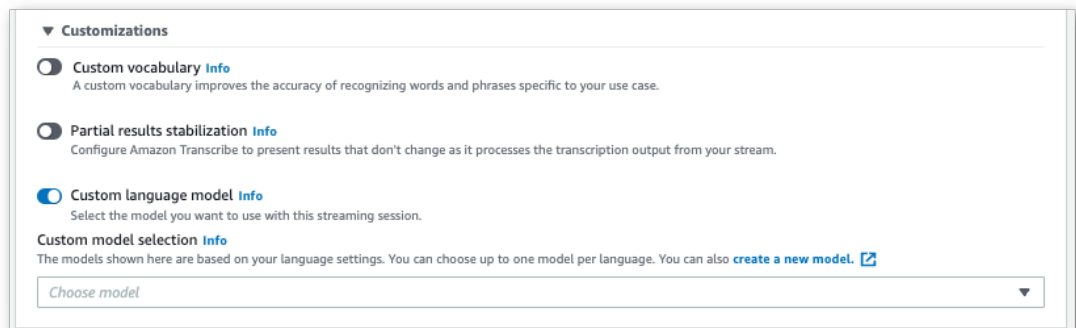
To use a custom language model with a streaming transcription, see the following for examples:

AWS Management Console

1. Sign into the [AWS Management Console](#).
2. In the navigation pane, choose **Real-time transcription**. Scroll down to **Customizations** and expand this field if it is minimized.



3. Toggle on **Custom language model** and select a model from the drop-down menu.



Include any other settings you wish to apply to your stream.

4. You're now ready to transcribe your stream. Select **Start streaming** and begin speaking. To end your dictation, select **Stop streaming**.

HTTP/2 stream

This example creates an HTTP/2 request that includes your custom language model. For more information on using HTTP/2 streaming with Amazon Transcribe, see [Setting up an HTTP/2 stream \(p. 82\)](#). For more detail on parameters and headers specific to Amazon Transcribe, see [StartStreamTranscription](#).

```
POST /stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
X-Amz-Target: com.amazonaws.transcribe.Transcribe.StartStreamTranscription
```

```
Content-Type: application/vnd.amazon.eventstream
X-Amz-Content-Sha256: string
X-Amz-Date: 20220208T235959Z
Authorization: AWS4-HMAC-SHA256 Credential=access-key/20220208/us-west-2/transcribe/
aws4_request, SignedHeaders=content-type;host;x-amz-content-sha256;x-amz-date;x-amz-
target;x-amz-security-token, Signature=string
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
x-amzn-transcribe-language-model-name: my-first-language-model
transfer-encoding: chunked
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

WebSocket stream

This example creates a pre-signed URL that applies your custom language model to a WebSocket stream. Line breaks have been added for readability. For more information on using WebSocket streams with Amazon Transcribe, see [Setting up a WebSocket stream \(p. 86\)](#). For more detail on parameters, see [StartStreamTranscription](#).

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/stream-transcription-websocket?
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=string
&X-Amz-SignedHeaders=content-type%3Bhost%3Bx-amz-date
&language-code=en-US
&media-encoding=flac
&sample-rate=16000
&language-model-name=my-first-language-model
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

Alternative transcriptions

When Amazon Transcribe transcribes audio, it creates different versions of the same transcript and assigns a confidence score to each version. In a typical transcription, you only get the version with the highest confidence score.

If you enable alternative transcriptions, Amazon Transcribe returns other versions of your transcript that have lower confidence levels. You can choose to have up to ten alternative transcriptions returned. If you specify a greater number of alternatives than what Amazon Transcribe identifies, only the actual number of alternatives is returned.

All alternatives are located in the same transcription output file and are presented at the segment level. Segments are natural pauses in speech, such as a change in speaker or a pause in the audio.

Alternative transcriptions are only available for batch transcriptions.

Your transcription output is structured as follows. The ellipses (...) in the code examples indicate where content has been removed for brevity.

1. A complete final transcription for a given segment.

```
"results": {
  "language_code": "en-US",
  "transcripts": [
    {
      "transcript": "The amazon is the largest rainforest on the planet."
    }
  ],
}
```

2. A confidence score for each word in the preceding transcript section.

```
"items": [
  {
    "start_time": "1.15",
    "end_time": "1.35",
    "alternatives": [
      {
        "confidence": "1.0",
        "content": "The"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "1.35",
    "end_time": "2.05",
    "alternatives": [
      {
        "confidence": "1.0",
        "content": "amazon"
      }
    ],
    "type": "pronunciation"
  },
]
```

3. Your alternative transcriptions are located in the `segments` portion of your transcription output. The alternatives for each segment are ordered by descending confidence score.

```
"segments": [
  {
    "start_time": "1.04",
    "end_time": "5.065",
    "alternatives": [
      {
        "transcript": "The amazon is the largest rain forest on the
planet.",
        "items": [
          {
            "start_time": "1.15",
            "confidence": "1.0",
            "end_time": "1.35",
            "type": "pronunciation",
            "content": "The"
          },
          {
            "start_time": "3.06",
            "confidence": "0.0037",
            "end_time": "3.38",
            "type": "pronunciation",
            "content": "rain"
          },
          {
            "start_time": "3.38",
            "confidence": "0.0037",
            "end_time": "3.96",
            "type": "pronunciation",
            "content": "forest"
          }
        ]
      }
    ]
  }
]
```

4. A status at the end of your transcription output.

```
"status": "COMPLETED"
}
```

Requesting alternative transcriptions

You can request alternative transcriptions using the [AWS Management Console](#), [AWS CLI](#), or [AWS SDK](#); see the following for examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This opens the **Specify job details** page.

Specify job details [Info](#)

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Model type [Info](#)

Choose the type of model to use for the transcription job.

General model
To use a model that is not specialized for a particular use case, choose this option. Configuration options vary between languages.

Custom language model
To use a model that you trained for your specific use case, choose this option. This model has fewer configuration options than the general model.

Language settings

You can transcribe your audio file in a language that you specify or have Amazon Transcribe Identify and transcribe it in the predominant language.

Specific language [Info](#)
If you know the language spoken in your source audio, choose this option to get the most accurate results. The options available for additional processing vary between languages.

Automatic language identification [Info](#)
If you don't know the language spoken in your audio files, choose this option. You have access to fewer options for additional processing than if you choose **Specific language**.

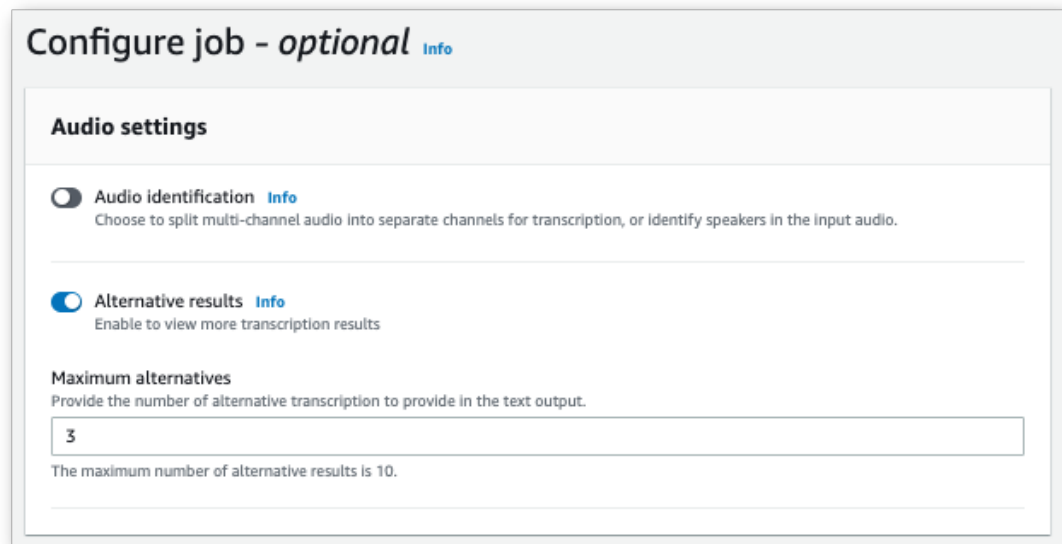
Language

Choose the language of the input audio.

[▶ Additional settings](#)

3. Fill in any fields you wish to include on the **Specify job details** page, then select **Next**. This takes you to the **Configure job - optional** page.

Select **Alternative results** and specify the maximum number of alternative transcription result you want in your transcript.



4. Select **Create job** to run your transcription job.

AWS CLI

This example uses the `start-transcription-job` command and `ShowAlternatives` parameter. For more information, see [StartTranscriptionJob](#) and [ShowAlternatives](#).

Note that if you include `ShowAlternatives=true` in your request, you must also include `MaxAlternatives`.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--language-code en-US \  
--settings ShowAlternatives=true,MaxAlternatives=4
```

Here's another example using the `start-transcription-job` command, and a request body that identifies the language.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://filepath/my-first-alt-transcription-job.json
```

The file `my-first-alt-transcription-job.json` contains the following request body.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "LanguageCode": "en-US",  
  "Settings": {
```

```
        "ShowAlternatives": true,  
        "MaxAlternatives": 4  
    }  
}
```

AWS SDK for Python (Boto3)

The following example uses the AWS SDK for Python (Boto3) to request alternative transcriptions by using the `ShowAlternatives` argument for the `start_transcription_job` method. For more information, see [StartTranscriptionJob](#) and [ShowAlternatives](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

Note that if you include `'ShowAlternatives': True` in your request, you must also include `MaxAlternatives`.

```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe', 'us-west-2')  
job_name = "my-first-transcription-job"  
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
transcribe.start_transcription_job(  
    TranscriptionJobName = job_name,  
    Media = {  
        'MediaFileUri': job_uri  
    },  
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',  
    OutputKey = 'my-output-files/',  
    LanguageCode = 'en-US',  
    Settings = {  
        'ShowAlternatives': True,  
        'MaxAlternatives': 4  
    }  
)  
  
while True:  
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)  
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:  
        break  
    print("Not ready yet...")  
    time.sleep(5)  
print(status)
```


Using custom vocabulary filters to delete, mask, or flag words

A custom vocabulary filter is a text file that contains a custom list of individual words that you want to modify in your transcription output.

A common use case is the removal of offensive or profane terms; however, custom vocabulary filters are completely custom, so you can select any words you want. For example, if you have a new product about to launch, you can mask the product name in meeting transcripts. In this case, you keep stakeholders up-to-date while keeping the product name secret until launch.

Vocabulary filtering has three display methods: `mask`, `remove`, and `tag`. Refer to the following examples to see how each works.

- **Mask:** Replaces specified words with three asterisks (***)

```
"transcript": "You can specify a list of *** or *** words, and *** *** removes them from transcripts automatically."
```

- **Remove:** Deletes specified words, leaving nothing in their place.

```
"transcript": "You can specify a list of or words, and removes them from transcripts automatically."
```

- **Tag:** Adds a tag (`"vocabularyFilterMatch": true`) to each specified word, but doesn't alter the word itself. Tagging allows for rapid transcript substitutions and edits.

```
"transcript": "You can specify a list of profane or offensive words, and amazon transcribe removes them from transcripts automatically."
...
  "alternatives": [
    {
      "confidence": "1.0",
      "content": "profane"
    }
  ],
  "type": "pronunciation",
  "vocabularyFilterMatch": true
```

When you submit a transcription request, you can specify a custom vocabulary filter and the filtration method you want to apply. Amazon Transcribe then modifies exact word matches when they appear in your transcript, according to the filtration method you specify.

Custom vocabulary filters can be applied to batch and streaming transcription requests. To learn how to create a custom vocabulary filter, see [Creating a vocabulary filter \(p. 132\)](#). To learn how to apply your custom vocabulary filter, see [Using a custom vocabulary filter \(p. 135\)](#).

For a video walkthrough of vocabulary filtering, see [Using vocabulary filters](#).

API operations specific to vocabulary filtering

[CreateVocabularyFilter](#), [DeleteVocabularyFilter](#), [GetVocabularyFilter](#), [ListVocabularyFilters](#), [UpdateVocabularyFilter](#)

Creating a vocabulary filter

There are two options for creating a custom vocabulary filter:

1. Save a list of line-separated words as a plain text file with UTF-8 encoding.
 - You can use this approach with the AWS Management Console, AWS CLI, or AWS SDKs.
 - If using the AWS Management Console, you can provide either a local path or an Amazon S3 URI for your custom vocabulary file.
 - If using the AWS CLI or AWS SDKs, you must upload your custom vocabulary file to an Amazon S3 bucket and include the Amazon S3 URI in your request.
2. Include a list of comma-separated words directly in your API request.
 - You can use this approach with the AWS CLI or AWS SDKs using the `words` parameter.

For examples of each method, refer to [Creating custom vocabulary filters \(p. 132\)](#)

Things to note when creating your custom vocabulary filter:

- Words aren't case sensitive. For example, "curse" and "CURSE" are treated the same.
- Only exact word matches are filtered. For example, if your filter includes "swear" but your media contains the word "swears" or "swearing", these are not filtered. Only instances of "swear" are filtered. You must therefore include all variations of the words you want filtered.
- Filters don't apply to words that are contained in other words. For example, if a custom vocabulary filter contains "marine" but not "submarine", "submarine" is not altered in the transcript.
- Each entry can only contain one word (no spaces).
- If you save your custom vocabulary filter as a text file, it must be in plain text format with UTF-8 encoding.
- You can have up to 100 custom vocabulary filters per AWS account and each can be up to 50 Kb in size.
- You can only use characters that are supported for your language. Refer to your language's [character set \(p. 5\)](#) for details.


Creating custom vocabulary filters

To process a custom vocabulary filter for use with Amazon Transcribe, see the following examples:

AWS Management Console

Before continuing, save your custom vocabulary filter as a text (*.txt) file. You can optionally upload your file to an Amazon S3 bucket.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Vocabulary filtering**. This opens the **Vocabulary filters** page where you can view existing custom vocabulary filters or create a new one.
3. Select **Create vocabulary filter**.



This takes you to the **Create vocabulary filter** page. Enter a name for your new custom vocabulary filter.

Select either the **File upload** or **S3 location** option under **Vocabulary input source**. Then specify the location of your custom vocabulary file.

Create vocabulary filter [Info](#)

Vocabulary filtering settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9 and - (hyphen).

Language

Vocabulary input source [Info](#)
 File upload
 S3 location

Vocabulary filter file location on S3
Provide a path to the S3 location where your vocabulary filter file is stored. To find a path, go to [Amazon S3](#) [↗](#)

File format: txt, maximum size 50 KB.

Tags - optional
A tag is a label you can add to a resource as metadata to help you organize, search, or filter your data. Each tag consists of a key and an optional value, in the form 'key:value'.
No tags associated with the resource.

You can add up to 50 more tags.

4. Optionally, add tags to your custom vocabulary filter. Once you have all fields completed, select **Create vocabulary filter** at the bottom of the page. If there are no errors processing your file, this takes you back to the **Vocabulary filters** page.

Your custom vocabulary filter is now ready to use.

AWS CLI

This example uses the `create-vocabulary-filter` command to process a word list into a usable custom vocabulary filter. For more information, see [CreateVocabularyFilter](#).

Option 1: You can include your list of words to your request using the `words` parameter.

```
aws transcribe create-vocabulary-filter \  
--vocabulary-filter-name my-first-vocabulary-filter \  
--language-code en-US \  
--words profane,offensive,Amazon,Transcribe
```

Option 2: You can save your list of words as a text file and upload it to an Amazon S3 bucket, then include the file's URI in your request using the `vocabulary-filter-file-uri` parameter.

```
aws transcribe create-vocabulary-filter \  
--vocabulary-filter-name my-first-vocabulary-filter \  
--language-code en-US \  
--vocabulary-filter-file-uri s3://DOC-EXAMPLE-BUCKET/my-vocabulary-filters/my-vocabulary-  
filter.txt
```

Here's another example using the [create-vocabulary-filter](#) command, and a request body that creates your custom vocabulary filter.

```
aws transcribe create-vocabulary-filter \  
--cli-input-json file://filepath/my-first-vocab-filter.json
```

The file *my-first-vocab-filter.json* contains the following request body.

Option 1: You can include your list of words to your request using the `words` parameter.

```
{  
  "VocabularyFilterName": "my-first-vocabulary-filter",  
  "LanguageCode": "en-US",  
  "Words": [  
    "profane", "offensive", "Amazon", "Transcribe"  
  ]  
}
```

Option 2: You can save your list of words as a text file and upload it to an Amazon S3 bucket, then include the file's URI in your request using the `VocabularyFilterFileUri` parameter.

```
{  
  "VocabularyFilterName": "my-first-vocabulary-filter",  
  "LanguageCode": "en-US",  
  "VocabularyFilterFileUri": "s3://DOC-EXAMPLE-BUCKET/my-vocabulary-filters/my-vocabulary-  
filter.txt"  
}
```

Note

If you include `VocabularyFilterFileUri` in your request, you cannot use `Words`; you must choose one or the other.

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to create a custom vocabulary filter using the [create_vocabulary_filter](#) method. For more information, see [CreateVocabularyFilter](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

Option 1: You can include your list of words to your request using the `words` parameter.

```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe', 'us-west-2')  
vocab_name = "my-first-vocabulary-filter"  
response = transcribe.create_vocabulary_filter(  
    LanguageCode = 'en-US',  
    VocabularyFilterName = vocab_name,  
    Words = [  

```

```
        'profane', 'offensive', 'Amazon', 'Transcribe'  
    ]  
)
```

Option 2: You can save your list of words as a text file and upload it to an Amazon S3 bucket, then include the file's URI in your request using the `VocabularyFilterFileUri` parameter.

```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe', 'us-west-2')  
vocab_name = "my-first-vocabulary-filter"  
response = transcribe.create_vocabulary_filter(  
    LanguageCode = 'en-US',  
    VocabularyFilterName = vocab_name,  
    VocabularyFilterFileUri = 's3://DOC-EXAMPLE-BUCKET/my-vocabulary-filters/my-vocabulary-  
filter.txt'  
)
```

Note

If you include `VocabularyFilterFileUri` in your request, you cannot use `Words`; you must choose one or the other.

Note

If you create a new Amazon S3 bucket for your custom vocabulary filter files, make sure the IAM role making the `CreateVocabularyFilter` request has permissions to access this bucket. If the role doesn't have the correct permissions, your request fails. You can optionally specify an IAM role within your request by including the `DataAccessRoleArn` parameter. For more information on IAM roles and policies in Amazon Transcribe, see [Amazon Transcribe identity-based policy examples \(p. 235\)](#).

Using a custom vocabulary filter

Once your custom vocabulary filter is created, you can include it in your transcription requests; refer to the following sections for examples.

The language of the custom vocabulary filter you're including in your request must match the language code you specify for your media. If you use language identification and specify multiple language options, you can include one custom vocabulary filter per specified language. If the languages of your custom vocabulary filters don't match the language identified in your audio, your filters are not applied to your transcription and there are no warnings or errors.

Using a custom vocabulary filter in a batch transcription

To use a custom vocabulary filter with a batch transcription, see the following for examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This opens the **Specify job details** page.

Specify job details [Info](#)

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Model type [Info](#)
Choose the type of model to use for the transcription job.

General model
To use a model that is not specialized for a particular use case, choose this option. Configuration options vary between languages.

Custom language model
To use a model that you trained for your specific use case, choose this option. This model has fewer configuration options than the general model.

Language settings
You can transcribe your audio file in a language that you specify or have Amazon Transcribe Identify and transcribe it in the predominant language.

Specific language [Info](#)
If you know the language spoken in your source audio, choose this option to get the most accurate results. The options available for additional processing vary between languages.

Automatic language identification [Info](#)
If you don't know the language spoken in your audio files, choose this option. You have access to fewer options for additional processing than if you choose **Specific language**.

Language
Choose the language of the input audio.

▶ **Additional settings**

- Name your job and specify your input media. Optionally include any other fields, then choose **Next**.
3. On the **Configure job** page, in the **Content removal** panel, toggle on **Vocabulary filtering**.

Configure job - *optional* [Info](#)

Audio settings

Audio identification [Info](#)
Choose to split multi-channel audio into separate channels for transcription, or identify speakers in the input audio.

Alternative results [Info](#)
Enable to view more transcription results

Content removal

Content removal conceals information in the resulting transcript from your source audio file. Amazon Transcribe changes items in the transcript and does not modify the source audio.

PII redaction [Info](#)
Label the type of PII and also mask the content with the PII entity type in the transcription output. For example, (123) 456-7890 will be masked as [PHONE].

Vocabulary filtering [Info](#)
Vocabulary filtering can remove, mask or tag specified words in the final transcript.

Filter selection
The vocabulary filters shown here are based on your language settings. You can choose up to one vocabulary filter per language. You can also [create a new vocabulary filter](#). [↗](#)

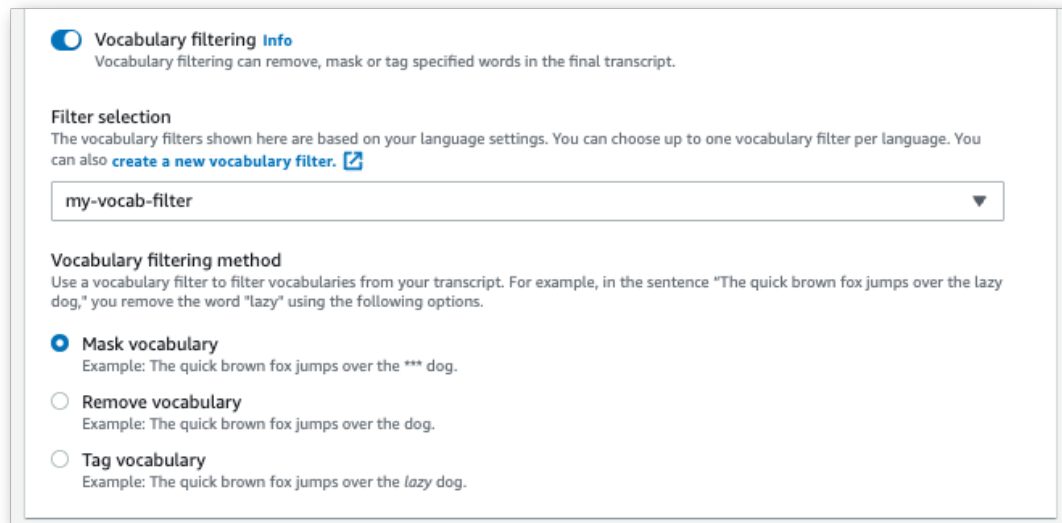
Choose a vocabulary filter ▼

Customization

Custom vocabulary [Info](#)
A custom vocabulary improves the accuracy of recognizing words and phrases specific to your use case.

Cancel

4. Select your custom vocabulary filter from the drop-down menu and specify the filtration method.



5. Select **Create job** to run your transcription job.

AWS CLI

This example uses the `start-transcription-job` command and `Settings` parameter with the `VocabularyFilterName` and `VocabularyFilterMethod` sub-parameters. For more information, see [StartTranscriptionJob](#) and [Settings](#).

```
aws transcribe start-transcription-job \
--region us-west-2 \
--transcription-job-name my-first-transcription-job \
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \
--output-bucket-name DOC-EXAMPLE-BUCKET \
--output-key my-output-files/ \
--language-code en-US \
--settings VocabularyFilterName=my-first-vocabulary-filter,VocabularyFilterMethod=mask
```

Here's another example using the `start-transcription-job` command, and a request body that includes your custom vocabulary with that job.

```
aws transcribe start-transcription-job \
--region us-west-2 \
--cli-input-json file://my-first-vocabulary-filter-job.json
```

The file `my-first-vocabulary-filter-job.json` contains the following request body.

```
{
  "TranscriptionJobName": "my-first-transcription-job",
  "Media": {
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
  },
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",
  "OutputKey": "my-output-files/",
  "LanguageCode": "en-US",
  "Settings": {
    "VocabularyFilterName": "my-first-vocabulary-filter",
    "VocabularyFilterMethod": "mask"
  }
}
```



```
}  
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to include a custom vocabulary filter using the `Settings` argument for the `start_transcription_job` method. For more information, see [StartTranscriptionJob](#) and [Settings](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

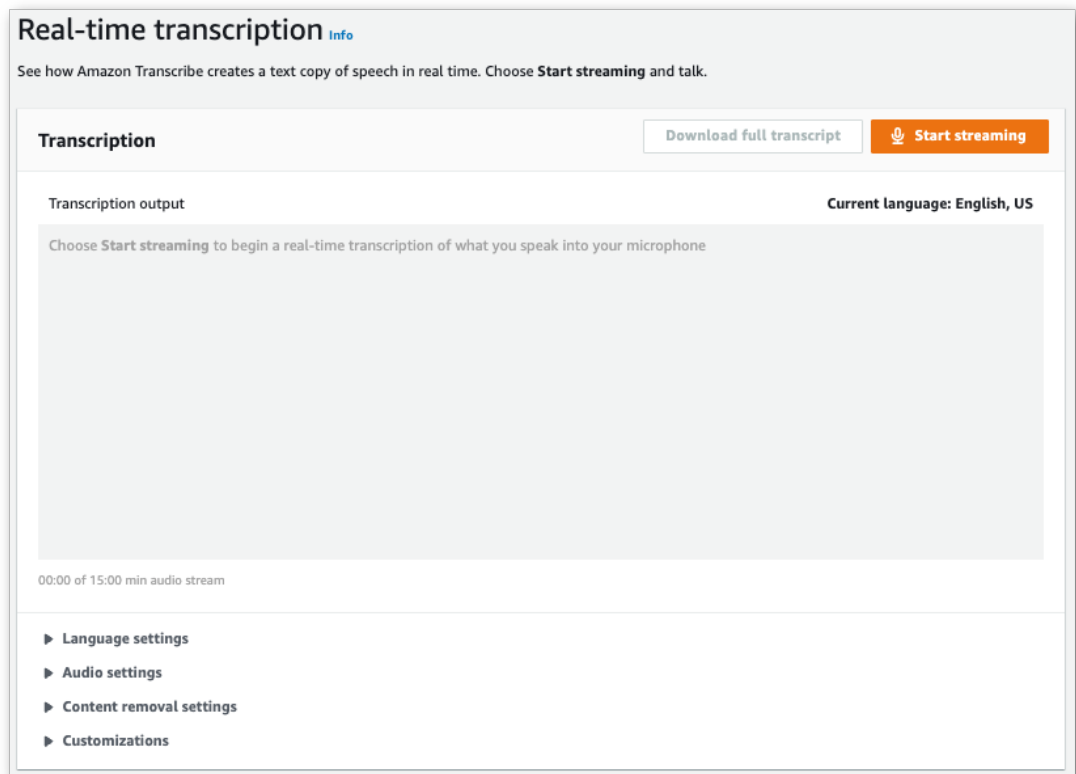
```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe', 'us-west-2')  
job_name = "my-first-transcription-job"  
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
transcribe.start_transcription_job(  
    TranscriptionJobName = job_name,  
    Media = {  
        'MediaFileUri': job_uri  
    },  
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',  
    OutputKey = 'my-output-files/',  
    LanguageCode = 'en-US',  
    Settings = {  
        'VocabularyFilterName': 'my-first-vocabulary-filter',  
        'VocabularyFilterMethod': 'mask'  
    }  
)  
  
while True:  
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)  
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:  
        break  
    print("Not ready yet...")  
    time.sleep(5)  
print(status)
```

Using a custom vocabulary filter in a streaming transcription

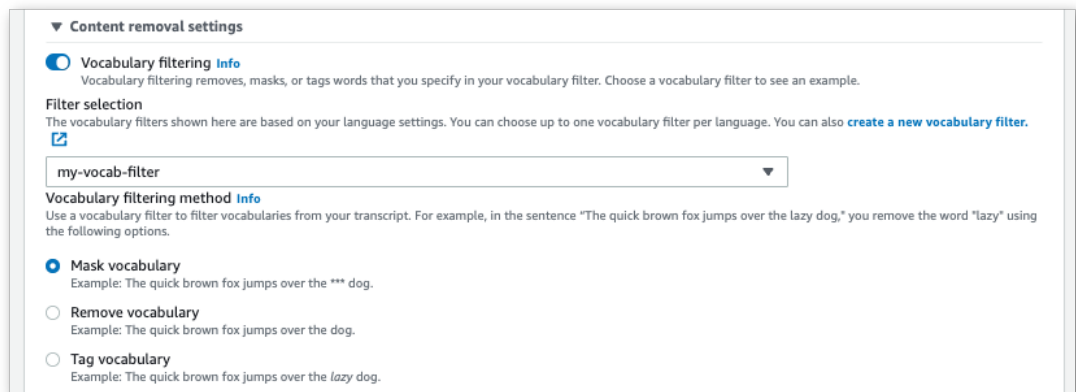
To use a custom vocabulary filter with a streaming transcription, see the following for examples:

AWS Management Console

1. Sign into the [AWS Management Console](#).
2. In the navigation pane, choose **Real-time transcription**. Scroll down to **Content removal settings** and expand this field if it is minimized.



3. Toggle on **Vocabulary filtering**. Select a custom vocabulary filter from the drop-down menu and specify the filtration method.



Include any other settings you wish to apply to your stream.

4. You're now ready to transcribe your stream. Select **Start streaming** and begin speaking. To end your dictation, select **Stop streaming**.

HTTP/2 stream

This example creates an HTTP/2 request that includes your custom vocabulary filter and filter method. For more information on using HTTP/2 streaming with Amazon Transcribe, see [Setting up an HTTP/2 stream \(p. 82\)](#). For more detail on parameters and headers specific to Amazon Transcribe, see [StartStreamTranscription](#).

Amazon Transcribe Developer Guide

Using a custom vocabulary filter in a streaming transcription

```
POST /stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
X-Amz-Target: com.amazonaws.transcribe.Transcribe.StartStreamTranscription
Content-Type: application/vnd.amazon.eventstream
X-Amz-Content-Sha256: string
X-Amz-Date: 20220208T235959Z
Authorization: AWS4-HMAC-SHA256 Credential=access-key/20220208/us-west-2/transcribe/
aws4_request, SignedHeaders=content-type;host;x-amz-content-sha256;x-amz-date;x-amz-
target;x-amz-security-token, Signature=string
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
x-amzn-transcribe-vocabulary-filter-name: my-first-vocabulary-filter
x-amzn-transcribe-vocabulary-filter-method: mask
transfer-encoding: chunked
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

WebSocket stream

This example creates a pre-signed URL that applies your custom vocabulary to a WebSocket stream. Line breaks have been added for readability. For more information on using WebSocket streams with Amazon Transcribe, see [Setting up a WebSocket stream \(p. 86\)](#). For more detail on parameters, see [StartStreamTranscription](#).

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/stream-transcription-websocket?
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=string
&X-Amz-SignedHeaders=content-type%3Bhost%3Bx-amz-date
&language-code=en-US
&media-encoding=flac
&sample-rate=16000
&vocabulary-filter-name=my-first-vocabulary-filter
&vocabulary-filter-method=mask
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

Redacting or identifying personally identifiable information

Redaction is used to mask or remove sensitive content, in the form of personally identifiable information (PII), from your transcripts. Amazon Transcribe can redact information with batch jobs and streaming transcriptions. Additionally, if you're performing a streaming transcription, you also have the option to flag PII without redacting it; see [Example PII identification output \(p. 152\)](#) for an example.

Types of PII Amazon Transcribe can recognize

PII type	Description
Bank Account Number	A number that uniquely identifies a bank account.
Bank Routing Number	A number that identifies the location of a bank account.
Credit Card Number or Debit Card Number	A value that uniquely defines a payment card issued by a bank.
Credit Card or Debit Card CVV Code	A 3-digit or 4-digit security code on each credit card.
Credit Card or Debit Card Expiration Date	The month and year a card expires.
Debit or credit card PIN	A security code issued by a bank or credit union. This number is used for bank accounts and payment cards.
Email address	The unique identifier of an email box where messages are delivered.
US mailing address	The United States mailing address of an individual.
Name	The first and last name of a person.
US phone number	A 10-digit phone number within the United States.
Social Security Number	A 9-digit number (or the last 4 digits of that number). Issued to US citizens, permanent residents, and temporary residents with employment.
All PII	Redact or identify all PII types listed in this table.

When redaction is enabled, you have the option to generate only a redacted transcript or both a redacted transcript and an unredacted transcript. If you choose to generate only a redacted transcript, note that your media is the only place where the complete conversation is stored. If you delete your original media, there is no record of the unredacted PII. Because of this, it may be prudent to generate an unredacted transcript in addition to a redacted one.

To learn more about PII redaction with batch transcriptions, refer to: [Redacting PII in your batch job \(p. 143\)](#).

To learn more about PII redaction or identification with streaming transcriptions, refer to: [Redacting or identifying PII in a real-time stream \(p. 146\)](#).

Important

The redaction feature is designed to identify and remove sensitive data. However, due to the predictive nature of machine learning, Amazon Transcribe may not identify and remove all instances of sensitive data in your transcript. We strongly recommend that you review any redacted output to ensure it meets your needs.

The redaction feature does not meet the requirements for de-identification under medical privacy laws, such as the U.S. Health Insurance Portability and Accountability Act of 1996 (HIPAA).

Redaction is supported with US English (en-US) only.

For a video walkthrough of the Amazon Transcribe's redaction feature, see [Use content redaction to identify & redact PII](#).

Redacting PII in your batch job

When redacting personally identifiable information (PII) from a transcript during a batch transcription job, Amazon Transcribe replaces each identified instance of PII with [PII] in the main text body of your transcript. You can also view the type of PII that was redacted in the word-for-word portion of the transcription output. For an output sample, see [Example redacted batch output \(p. 149\)](#).

Both redacted and unredacted transcripts are stored in the same output Amazon S3 bucket. Amazon Transcribe stores them in either a bucket you specify or in the default Amazon S3 bucket managed by the service.

You can start a batch transcription job using the AWS Management Console, AWS CLI, or AWS SDK.

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This will open the **Specify job details** page.
3. After filling in your desired fields on the **Specify job details** page, select **Next** to go to the **Configure job - optional** page. Here you'll find the **Content removal** panel with the **PII redaction** toggle.

Configure job - *optional* [Info](#)

Audio settings

Audio identification [Info](#)
Choose to split multi-channel audio into separate channels for transcription, or identify speakers in the input audio.

Alternative results [Info](#)
Enable to view more transcription results

Content removal

Content removal conceals information in the resulting transcript from your source audio file. Amazon Transcribe changes items in the transcript and does not modify the source audio.

PII redaction [Info](#)
Label the type of PII and also mask the content with the PII entity type in the transcription output. For example, (123) 456-7890 will be masked as [PHONE].

4. Once you select **PII redaction**, you have the option to select all PII types you want to redact. You can also choose to have an unredacted transcript if you select **Include unredacted transcript in job output** box.

Content removal

Content removal conceals information in the resulting transcript from your source audio file. Amazon Transcribe changes items in the transcript and does not modify the source audio.

PII redaction [Info](#)
Label the type of PII and also mask the content with the PII entity type in the transcription output. For example, (123) 456-7890 will be masked as [PHONE].

Include unredacted transcript in job output
Returns unredacted version of the transcript in addition to the redacted version.

Select PII entity types (11 of 11 selected)

Select All

Financial (6 of 6 selected)

<input checked="" type="checkbox"/> BANK_ACCOUNT_NUMBER	<input checked="" type="checkbox"/> BANK_ROUTING	<input checked="" type="checkbox"/> CREDIT_DEBIT_NUMBER
<input checked="" type="checkbox"/> CREDIT_DEBIT_CVV	<input checked="" type="checkbox"/> CREDIT_DEBIT_EXPIRY	<input checked="" type="checkbox"/> PIN

Personal (5 of 5 selected)

<input checked="" type="checkbox"/> NAME	<input checked="" type="checkbox"/> ADDRESS	<input checked="" type="checkbox"/> PHONE
<input checked="" type="checkbox"/> EMAIL	<input checked="" type="checkbox"/> SSN	

Vocabulary filtering [Info](#)
Vocabulary filtering can remove, mask or tag specified words in the final transcript.

5. Select **Create job** to run your transcription job.

AWS CLI

This example uses the `start-transcription-job` command and `content-redaction` parameter. For more information, see [StartTranscriptionJob](#) and [ContentRedaction](#).

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--language-code en-US \  
--content-redaction  
RedactionType=PII,RedactionOutput=redacted,PiiEntityTypes=NAME,ADDRESS,BANK_ACCOUNT_NUMBER
```

Here's another example using the `start-transcription-job` method, and the request body redacts PII for that job.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://filepath/my-first-redaction-job.json
```

The file `my-first-redaction-job.json` contains the following request body.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "LanguageCode": "en-US",  
  "ContentRedaction": {  
    "RedactionOutput": "redacted",  
    "RedactionType": "PII",  
    "PiiEntityTypes": [  
      "NAME",  
      "ADDRESS",  
      "BANK_ACCOUNT_NUMBER"  
    ]  
  }  
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to redact content using the `ContentRedaction` argument for the `start_transcription_job` method. For more information, see [StartTranscriptionJob](#) and [ContentRedaction](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function  
import time  
import boto3
```

```
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
transcribe.start_transcription_job(
    TranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'my-output-files/',
    LanguageCode = 'en-US',
    ContentRedaction = {
        'RedactionOutput': 'redacted',
        'RedactionType': 'PII',
        'PiiEntityTypes': [
            'NAME', 'ADDRESS', 'BANK_ACCOUNT_NUMBER'
        ]
    }
)

while True:
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Note

PII redaction for batch jobs is only supported in these AWS Regions: Asia Pacific (Hong Kong), Asia Pacific (Mumbai), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), GovCloud (US-West), Canada (Central), EU (Frankfurt), EU (Ireland), EU (London), EU (Paris), Middle East (Bahrain), South America (Sao Paulo), US East (N. Virginia), US East (Ohio), US West (Oregon), and US West (N. California).

Redacting or identifying PII in a real-time stream

When redacting personally identifiable information (PII) from a streaming transcription, Amazon Transcribe replaces each identified instance of PII with [PII] in your transcript.

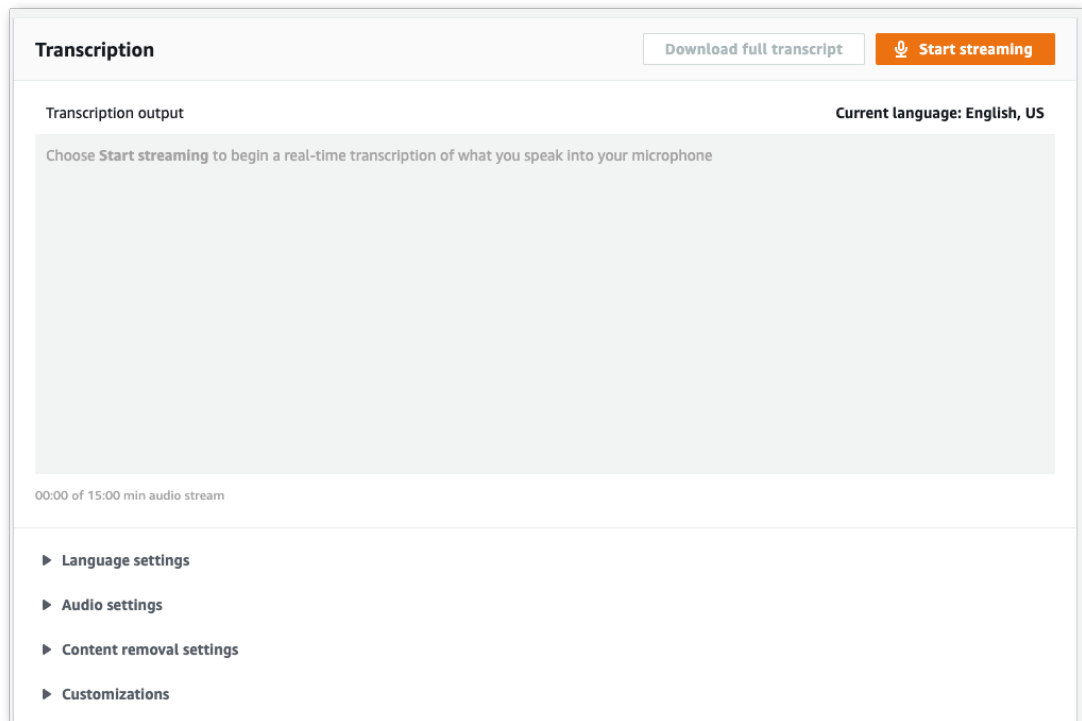
An additional option available for streaming audio is *PII identification*. When you activate PII Identification, Amazon Transcribe labels the PII in your transcription results under an `Entities` object. For an output sample, see [Example redacted streaming output \(p. 151\)](#) and [Example PII identification output \(p. 152\)](#).

PII identification and redaction for streaming jobs is performed only upon complete transcription of the audio segments.


You can start a streaming transcription using the AWS Management Console, WebSocket, or HTTP/2.

AWS Management Console

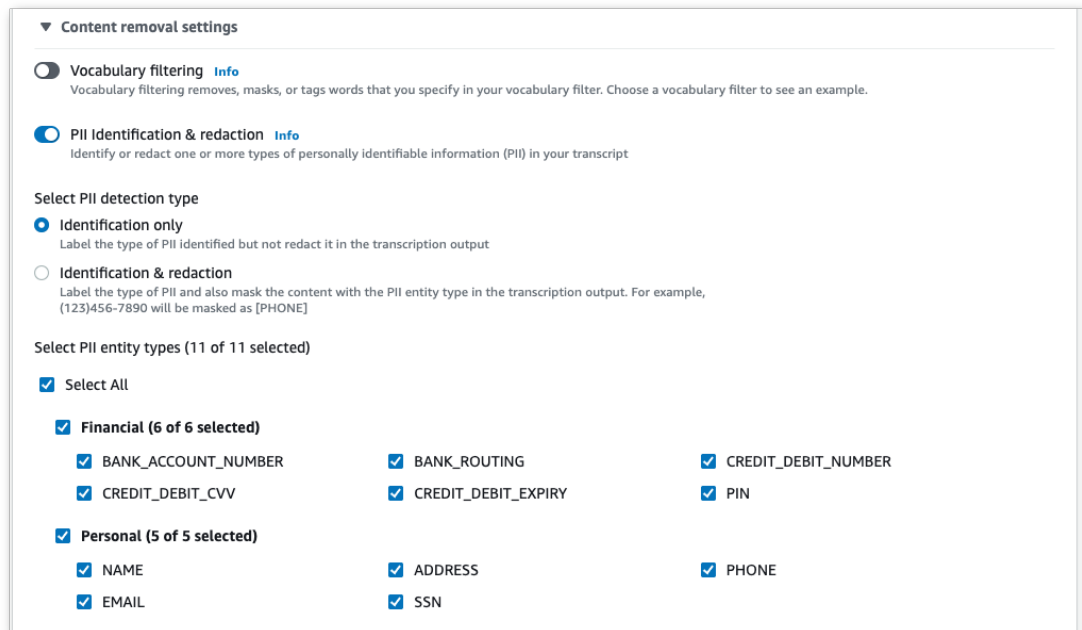
1. Sign into the [AWS Management Console](#).
2. In the navigation pane, choose **Real-time transcription**. Scroll down to **Content removal settings** and expand this field if it is minimized.



3. Toggle on **PII Identification & redaction**.



4. Select either **Identification only** or **Identification & redaction**, then select the PII entity types you want to identify or redact in your transcript.



Content removal settings

Vocabulary filtering [Info](#)
Vocabulary filtering removes, masks, or tags words that you specify in your vocabulary filter. Choose a vocabulary filter to see an example.

PII Identification & redaction [Info](#)
Identify or redact one or more types of personally identifiable information (PII) in your transcript

Select PII detection type

Identification only
Label the type of PII identified but not redact it in the transcription output

Identification & redaction
Label the type of PII and also mask the content with the PII entity type in the transcription output. For example, (123)456-7890 will be masked as [PHONE]

Select PII entity types (11 of 11 selected)

Select All

Financial (6 of 6 selected)

<input checked="" type="checkbox"/> BANK_ACCOUNT_NUMBER	<input checked="" type="checkbox"/> BANK_ROUTING	<input checked="" type="checkbox"/> CREDIT_DEBIT_NUMBER
<input checked="" type="checkbox"/> CREDIT_DEBIT_CVV	<input checked="" type="checkbox"/> CREDIT_DEBIT_EXPIRY	<input checked="" type="checkbox"/> PIN

Personal (5 of 5 selected)

<input checked="" type="checkbox"/> NAME	<input checked="" type="checkbox"/> ADDRESS	<input checked="" type="checkbox"/> PHONE
<input checked="" type="checkbox"/> EMAIL	<input checked="" type="checkbox"/> SSN	

5. You're now ready to transcribe your stream. Select **Start streaming** and begin speaking. To end your dictation, select **Stop streaming**.

WebSocket stream

This example creates a pre-signed URL that uses either PII Identification or PII redaction in a WebSocket stream. Line breaks have been added for readability. For more information on using WebSocket streams with Amazon Transcribe, see [Setting up a WebSocket stream \(p. 86\)](#). For more detail on parameters, see [StartStreamTranscription](#).

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/stream-transcription-websocket?
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=string
&X-Amz-SignedHeaders=content-type%3Bhost%3Bx-amz-date
&language-code=en-US
&media-encoding=flac
&sample-rate=16000
&pii-entity-types=NAME,ADDRESS
&content-redaction-type=PII (or &content-identification-type=PII)
```

You cannot use both `content-identification-type` and `content-redaction-type` in the same request.

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

HTTP/2 stream

This example creates an HTTP/2 request with PII identification or PII redaction enabled. For more information on using HTTP/2 streaming with Amazon Transcribe, see [Setting up an HTTP/2](#)

[stream \(p. 82\)](#). For more detail on parameters and headers specific to Amazon Transcribe, see [StartStreamTranscription](#).

```
POST /stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
X-Amz-Target: com.amazonaws.transcribe.Transcribe.StartStreamTranscription
Content-Type: application/vnd.amazon.eventstream
X-Amz-Content-Sha256: string
X-Amz-Date: 20220208T235959Z
Authorization: AWS4-HMAC-SHA256 Credential=access-key/20220208/us-west-2/transcribe/
aws4_request, SignedHeaders=content-type;host;x-amz-content-sha256;x-amz-date;x-amz-
target;x-amz-security-token, Signature=string
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
x-amzn-transcribe-content-identification-type: PII (or x-amzn-transcribe-content-redaction-
type: PII)
x-amzn-transcribe-pii-entity-types: NAME,ADDRESS
transfer-encoding: chunked
```

You cannot use both `content-identification-type` and `content-redaction-type` in the same request.

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

Note

PII redaction for streaming is only supported in these AWS Regions: Asia Pacific (Seoul), Asia Pacific (Sydney), Asia Pacific (Tokyo), Canada (Central), EU (Frankfurt), EU (Ireland), EU (London), US East (N. Virginia), US East (Ohio), and US West (Oregon).

Example PII redaction and identification output

The following examples show redacted output from batch and streaming jobs, and PII identification from a streaming job.

Transcription jobs using content redaction generate two types of confidence values. The Automatic Speech Recognition (ASR) confidence indicates the items that have the type of pronunciation or punctuation is a specific utterance. In the following transcript output, the word `Good` has a confidence of `1.0`. This confidence value indicates that Amazon Transcribe is 100 percent confident that the word uttered in this transcript is 'Good'. The confidence value for a `[PII]` tag is the confidence that the speech it flagged for redaction is truly PII. In the following transcript output, the confidence of `0.9999` indicates that Amazon Transcribe is 99.99 percent confident that the entity it redacted in the transcript is PII.

Example redacted batch output

```
{
  "jobName": "my-first-transcription-job",
  "accountId": "111122223333",
  "isRedacted": true,
  "results": {
    "transcripts": [
      {
        "transcript": "Good morning, everybody. My name is [PII], and today I feel
like
Social          sharing a whole lot of personal information with you. Let's start with my
```

```
[PII]. Security number [PII]. My credit card number is [PII] and my C V V code is
personal I hope that Amazon Transcribe is doing a good job at redacting that
information away. Let's check."
}
],
"items": [
  {
    "start_time": "2.86",
    "end_time": "3.35",
    "alternatives": [
      {
        "confidence": "1.0",
        "content": "Good"
      }
    ],
    "type": "pronunciation"
  },
  Items removed for brevity
  {
    "start_time": "5.56",
    "end_time": "6.25",
    "alternatives": [
      {
        "content": "[PII]",
        "redactions": [
          {
            "confidence": "0.9999",
            "type": "NAME",
            "category": "PII"
          }
        ]
      }
    ],
    "type": "pronunciation"
  },
  Items removed for brevity
],
"status": "COMPLETED"
}
```

Here's the unredacted transcript for comparison:

```
{
  "jobName": "job id",
  "accountId": "111122223333",
  "isRedacted": false,
  "results": {
    "transcripts": [
      {
        "transcript": "Good morning, everybody. My name is Mike, and today I feel
like
Social
job
sharing a whole lot of personal information with you. Let's start with my
Security number 00000000. My credit card number is 5555555555555555
and my C V V code is 000. I hope that Amazon Transcribe is doing a good
at redacting that personal information away. Let's check."
      }
    ],
    "items": [
      {
        "start_time": "2.86",
```

```
        "end_time": "3.35",
        "alternatives": [
          {
            "confidence": "1.0",
            "content": "Good"
          }
        ],
        "type": "pronunciation"
      },
      Items removed for brevity
    {
      "start_time": "5.56",
      "end_time": "6.25",
      "alternatives": [
        {
          "confidence": "0.9999",
          "content": "Mike",
        }
      ],
      "type": "pronunciation"
    },
    Items removed for brevity
  ],
},
"status": "COMPLETED"
}
```

Example redacted streaming output

```
{
  "TranscriptResultStream": {
    "TranscriptEvent": {
      "Transcript": {
        "Results": [
          {
            "Alternatives": [
              {
                "Transcript": "my name is [NAME]",
                "Items": [
                  {
                    "Content": "my",
                    "EndTime": 0.3799375,
                    "StartTime": 0.0299375,
                    "Type": "pronunciation"
                  },
                  {
                    "Content": "name",
                    "EndTime": 0.5899375,
                    "StartTime": 0.3899375,
                    "Type": "pronunciation"
                  },
                  {
                    "Content": "is",
                    "EndTime": 0.7899375,
                    "StartTime": 0.5999375,
                    "Type": "pronunciation"
                  },
                  {
                    "Content": "[NAME]",
                    "EndTime": 1.0199375,
                    "StartTime": 0.7999375,
                    "Type": "pronunciation"
                  }
                ]
              }
            ]
          }
        ]
      }
    }
  }
}
```

```
    ],
    "Entities": [
      {
        "Content": "[NAME]",
        "Category": "PII",
        "Type": "NAME",
        "StartTime": 0.7999375,
        "EndTime": 1.0199375,
        "Confidence": 0.9989
      }
    ]
  },
  "EndTime": 1.02,
  "IsPartial": false,
  "ResultId": "2db76dc8-d728-11e8-9f8b-f2801f1b9fd1",
  "StartTime": 0.0199375
}
]
}
}
```

Example PII identification output

PII identification is an additional feature that you can use with your streaming transcription job.

```
{
  "TranscriptResultStream": {
    "TranscriptEvent": {
      "Transcript": {
        "Results": [
          {
            "Alternatives": [
              {
                "Transcript": "my name is mike",
                "Items": [
                  {
                    "Content": "my",
                    "EndTime": 0.3799375,
                    "StartTime": 0.0299375,
                    "Type": "pronunciation"
                  },
                  {
                    "Content": "name",
                    "EndTime": 0.5899375,
                    "StartTime": 0.3899375,
                    "Type": "pronunciation"
                  },
                  {
                    "Content": "is",
                    "EndTime": 0.7899375,
                    "StartTime": 0.5999375,
                    "Type": "pronunciation"
                  },
                  {
                    "Content": "mike",
                    "EndTime": 0.9199375,
                    "StartTime": 0.7999375,
                    "Type": "pronunciation"
                  }
                ]
              }
            ]
          }
        ]
      }
    }
  }
}
```

```
    ],
    "Entities": [
      {
        "Content": "mike",
        "Category": "PII",
        "Type": "NAME",
        "StartTime" : 0.7999375,
        "EndTime" : 1.0199375,
        "Confidence": 0.9989
      }
    ]
  },
  "EndTime": 1.02,
  "IsPartial": false,
  "ResultId": "2db76dc8-d728-11e8-9f8b-f2801f1b9fd1",
  "StartTime": 0.0199375
}
}
}
}
```

Transcribing multi-channel audio

If your audio has multiple channels, you can use *channel identification* to transcribe the speech from each of those channels. Amazon Transcribe identifies the speech from each channel and transcribes that speech in separate transcriptions. It combines those transcriptions into a single transcription output.

You can enable channel identification for both batch processing and real-time streaming. The following list describes how to enable it for each method.

- Batch transcription - [AWS Management Console](#) and [StartTranscriptionJob](#) API
- Streaming transcription - WebSocket streaming and [StartStreamTranscription](#) API

Transcribing multi-channel audio

To transcribe multi-channel audio in a batch transcription job, use the [AWS Management Console](#) or the [StartTranscriptionJob](#) API.

AWS Management Console

To use the AWS Management Console to enable channel identification in your batch transcription job, you enable audio identification and then channel identification. Channel identification is a subset of audio identification in the AWS Management Console.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, provide information about your transcription job.
5. Choose **Next**.
6. Enable **Audio identification**.
7. For **Audio identification type**, choose **Channel identification**.
8. Choose **Create**.

API

To transcribe multi-channel audio (API)

- For the [StartTranscriptionJob](#) API, specify the following.
 - a. For `TranscriptionJobName`, specify a name unique to your AWS account.
 - b. For `LanguageCode`, specify the language code that corresponds to the language spoken in your media file. For available languages and corresponding language codes, see [What is Amazon Transcribe? \(p. 1\)](#).
 - c. For the `MediaFileUri` parameter of the `Media` object, specify the name of the media file you want to transcribe.
 - d. For the `Settings` object, set `ChannelIdentification` to `true`.

The following is an example request using the AWS SDK for Python (Boto3).

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe')
job_name = "my-first-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
transcribe.start_transcription_job(
    TranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    LanguageCode = 'en-US',
    Settings = {
        'ChannelIdentification': True,
    }
)

while True:
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

AWS CLI

- Run the following code.

```
aws transcribe start-transcription-job \
--cli-input-json file://filepath/my-first-channel-id-job.json
```

The following is the code of `my-first-channel-id-job.json`.

```
{
  "TranscriptionJobName": "my-first-transcription-job",
  "LanguageCode": "en-US",
  "Media": {
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
  },
  "Settings": {
    "ChannelIdentification": true
  }
}
```

The following code shows the transcription output for audio that has a conversation on two channels.

```
{
  "jobName": "my-first-transcription-job",
  "accountId": "111122223333",
  "results": {
    "transcripts": [
      {
```

```
    "transcript": "When you try ... It seems to ..."  
  }  
],  
"channel_labels": {  
  "channels": [  
    {  
      "channel_label": "ch_0",  
      "items": [  
        {  
          "start_time": "12.282",  
          "end_time": "12.592",  
          "alternatives": [  
            {  
              "confidence": "1.0000",  
              "content": "When"  
            }  
          ],  
          "type": "pronunciation"  
        },  
        {  
          "start_time": "12.592",  
          "end_time": "12.692",  
          "alternatives": [  
            {  
              "confidence": "0.8787",  
              "content": "you"  
            }  
          ],  
          "type": "pronunciation"  
        },  
        {  
          "start_time": "12.702",  
          "end_time": "13.252",  
          "alternatives": [  
            {  
              "confidence": "0.8318",  
              "content": "try"  
            }  
          ],  
          "type": "pronunciation"  
        },  
        ...  
      ]  
    },  
    {  
      "channel_label": "ch_1",  
      "items": [  
        {  
          "start_time": "12.379",  
          "end_time": "12.589",  
          "alternatives": [  
            {  
              "confidence": "0.5645",  
              "content": "It"  
            }  
          ],  
          "type": "pronunciation"  
        },  
        {  
          "start_time": "12.599",  
          "end_time": "12.659",  
          "alternatives": [  
            {  
              "confidence": "0.2907",  
              "content": "seems"  
            }  
          ]  
        }  
      ]  
    }  
  ]  
}
```

```
    ],  
    "type": "pronunciation"  
  },  
  {  
    "start_time": "12.669",  
    "end_time": "13.029",  
    "alternatives": [  
      {  
        "confidence": "0.2497",  
        "content": "to"  
      }  
    ],  
    "type": "pronunciation"  
  },  
  ...  
]  
}
```

Amazon Transcribe transcribes the audio from each channel separately and combines the transcribed text from each channel into a single transcription output.

For each channel in the transcription output, Amazon Transcribe returns a list of *items*. An item is a transcribed word, pause, or punctuation mark. Each item has a start time and an end time. If a person on one channel speaks over a person on a separate channel, the start times and end times of the items for each channel overlap while the individuals are speaking over each other.

Transcribing multi-channel audio streams

You can transcribe audio from separate channels in either HTTP/2 or WebSocket streams using the [StartStreamTranscription](#) API.

Transcribing multi-channel audio in an HTTP/2 stream

To transcribe multi-channel audio in an HTTP/2 stream, use the [StartStreamTranscription](#) API and specify the following:

- `LanguageCode` - The language code of the audio.
- `MediaEncoding` - The encoding of the audio.
- `EnableChannelIdentification` - `true`
- `NumberOfChannels` - The number of channels in your streaming audio.

The following is the syntax for the parameters of an HTTP/2 request.

```
POST /stream-transcription HTTP/2  
host: transcribestreaming.us-west-2.amazonaws.com  
authorization: Generated value  
x-amz-target: com.amazonaws.transcribe.Transcribe.StartStreamTranscription  
x-amz-content-sha256: STREAMING-AWS4-HMAC-SHA256-EVENTS  
x-amz-date: 20220208T235959Z  
x-amzn-transcribe-session-id: my-first-http2-channel-id-stream  
x-amzn-transcribe-language-code: en-US  
x-amzn-transcribe-media-encoding: flac  
x-amzn-transcribe-sample-rate: 16000  
x-amzn-transcribe-enable-channel-identification: true
```

```
x-amzn-transcribe-number-of-channels: 2
content-type: application/vnd.amazon.eventstream
transfer-encoding: chunked
```

Parameter descriptions:

- **host:** Update the AWS Region ('us-west-2' in the preceding example) with the AWS Region you are calling. For a list of valid AWS Regions, see [AWS Regions and Endpoints](#).
- **authorization:** This is a generated field. To learn more about creating a signature, see [Signing AWS requests with Signature Version 4](#).
- **x-amz-target:** Don't alter this field; use the content shown in the preceding example.
- **x-amz-content-sha256:** This is a generated field. To learn more about calculating a signature, see [Signing AWS requests with Signature Version 4](#).
- **x-amz-date:** The date and time the signature is created. The format is YYYYMMDDTHHMMSSZ, where YYYY=year, MM=month, DD=day, HH=hour, MM=minute, SS=seconds, and 'T' and 'Z' are fixed characters. For more information, refer to [Handling Dates in Signature Version 4](#).
- **x-amzn-transcribe-session-id:** The name for your streaming session.
- **x-amzn-transcribe-language-code:** The encoding used for your input audio. Refer to [StartStreamTranscription](#) or [Supported languages and language-specific features \(p. 3\)](#) for a list of valid values.
- **x-amzn-transcribe-media-encoding:** The encoding used for your input audio. Valid values are pcm, ogg-opus, and flac.
- **x-amzn-transcribe-sample-rate:** The sample rate of the input audio (in Hertz). Amazon Transcribe supports a range from 8,000 Hz to 48,000 Hz. Low-quality audio, such as telephone audio, is typically around 8,000 Hz. High-quality audio typically ranges from 16,000 Hz to 48,000 Hz. Note that the sample rate you specify **must** match that of your audio.
- **x-amzn-transcribe-enable-channel-identification:** To enable channel identification, this field must be true.
- **x-amzn-transcribe-number-of-channels:** Must be 2 when x-amzn-transcribe-enable-channel-identification is set to true.
- **content-type:** Don't alter this field; use the content shown in the preceding example.

Transcribing multi-channel audio in a WebSocket stream

To identify speakers in WebSocket streams, use the following format to create a pre-signed URI and start a WebSocket request. Specify `enable-channel-id` as `true` and the number of channels in your stream in `number-of-channels`. A pre-signed URI contains the information needed to set up bi-directional communication between your application and Amazon Transcribe.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/stream-transcription-websocket
?language-code=LanguageCode
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=Signature Version 4 signature
&X-Amz-SignedHeaders=host
&media-encoding=flac
&sample-rate=16000
&session-id=sessionId
&enable-channel-identification=true
```

```
&number-of-channels=2
```

Multi-channel streaming output

The output of a streaming transcription is the same for HTTP/2 and WebSocket requests. The following is an example output.

```
{
  "resultId": "XXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXX",
  "startTime": 0.11,
  "endTime": 0.66,
  "isPartial": false,
  "alternatives": [
    {
      "transcript": "Left.",
      "items": [
        {
          "startTime": 0.11,
          "endTime": 0.45,
          "type": "pronunciation",
          "content": "Left",
          "vocabularyFilterMatch": false
        },
        {
          "startTime": 0.45,
          "endTime": 0.45,
          "type": "punctuation",
          "content": ".",
          "vocabularyFilterMatch": false
        }
      ]
    }
  ],
  "channelId": "ch_0"
}
```

For each speech segment, there is a `channelId` flag that indicates which channel the speech belongs to.

Identifying the dominant languages in your media

Amazon Transcribe is able to automatically identify the languages spoken in your media without you having to specify a language code.

[Batch language identification \(p. 160\)](#) can identify the dominant language spoken in your media file or, if your media contains multiple languages, it can identify all languages spoken. To improve language identification accuracy, you can optionally provide a list of two or more languages you think may be present in your media.

[Streaming language identification \(p. 167\)](#) can identify one language per channel (a maximum of two channels are supported). Streaming requests must have a minimum of two additional language options included in your request. Providing language options allows for faster language identification. The faster Amazon Transcribe is able to identify the language, the less change there is of data loss in the first few seconds of your stream.

Important

Batch and streaming transcriptions support different languages. Refer to the [supported languages table \(p. 3\)](#) for details.

To learn about monitoring and events with language identification, refer to [Language identification events \(p. 250\)](#).

Language identification with batch transcription jobs

Use batch language identification to automatically identify the language, or languages, in your media file.

If your media contains only one language, you can enable [single-language identification](#), which identifies the dominant language spoken in your media file and creates your transcript using only this language.

If your media contains more than one language, you can enable [multi-language identification](#), which identifies all languages spoken in your media file and creates your transcript using each identified language. Note that a multi-lingual transcript is produced. You can use other services, such as Amazon Translate, to translate your transcript.

Refer to the [supported languages \(p. 3\)](#) table for a complete list of supported languages and associated language codes. For best results, ensure that your media file contains at least 30 seconds of speech.

For usage examples with the AWS Management Console, AWS CLI, and AWS Python SDK, see [Using language identification with batch transcriptions \(p. 162\)](#).

Identifying languages in multi-language audio

Multi-language identification is intended for multi-lingual media files, and provides you with a transcript that reflects all [supported languages \(p. 3\)](#) spoken in your media. This means that if speakers change

languages mid-conversation, or if each participant is speaking a different language, your transcription output detects and transcribes each language correctly. For example, if your media contains a bilingual speaker who is alternating between US English (`en-US`) and Hindi (`hi-IN`), multi-language identification can identify and transcribe spoken US English as `en-US` and spoken Hindi as `hi-IN`.

This differs from single-language identification, where only one dominant language is used to create a transcript. In this case, any spoken language that is not the dominant language is transcribed incorrectly.

Note

Redaction and custom language models are not currently supported with multi-language identification.

Multi-language transcripts provide a summary of detected languages and the total time that each language is spoken in your media. Here's an example:

```
"results": {
  "transcripts": [
    {
      "transcript": "welcome to Amazon transcribe. ## ## ##### ### #### #### ##
##### ## #####"
    }
  ],
  ...
  "language_codes": [
    {
      "language_code": "en-US",
      "duration_in_seconds": 2.45
    },
    {
      "language_code": "hi-IN",
      "duration_in_seconds": 5.325
    },
    {
      "language_code": "ja-JP",
      "duration_in_seconds": 4.15
    }
  ]
}
```

Improving language identification accuracy

With language identification, you have the option to include a list of languages you think may be present in your media. Including language options (`LanguageOptions`) restricts Amazon Transcribe to using only the languages that you specify when matching your audio to the correct language, which can speed up language identification and improve the accuracy associated with assigning the correct language dialect.

If you choose to include language codes, you must include at least two. There's no limit on the number of language codes you can include, but we recommend using between two and five for optimal efficiency and accuracy.

Note

If you include language codes with your request and none of the language codes you provide match the language, or languages, identified in your audio, Amazon Transcribe selects the closest language match from your specified language codes. It then produces a transcript in that language. For example, if your media is in US English (`en-US`) and you provide Amazon Transcribe with the language codes `zh-CN`, `fr-FR`, and `de-DE`, Amazon Transcribe is likely to match your media to German (`de-DE`) and produce a German-language transcription.

Mismatching language codes and spoken languages can result in an inaccurate transcript, so we advise caution when including language codes.

Combining language identification with other Amazon Transcribe features

You can use batch language identification in combination with any other Amazon Transcribe feature. If combining language identification with other features, you are limited to the languages supported with those features. For example, if using language identification with content redaction, you are limited to US English (en-US), as this is only language available for redaction. Refer to [Supported languages and language-specific features](#) (p. 3) for more information.

Important

If you're using automatic language identification with content redaction enabled and your audio contains languages other than US English (en-US), only the US English content is redacted in your transcript. Other languages cannot be redacted and there are no warnings or job failures.

Custom language models, custom vocabularies, and custom vocabulary filters

If you want to add one or more custom language models, custom vocabularies, or custom vocabulary filters to your language identification request, you must include the [LanguageIdSettings](#) parameter. You can then specify a language code with a corresponding custom language model, custom vocabulary, and custom vocabulary filter. Note that multi-language identification doesn't support custom language models.

It's recommended that you include `LanguageOptions` when using [LanguageIdSettings](#) to ensure that the correct language dialect is identified. For example, if you specify an en-US custom vocabulary, but Amazon Transcribe determines that the language spoken in your media is en-AU, your custom vocabulary *is not* applied to your transcription. If you include `LanguageOptions` and specify en-US as the only English language dialect, your custom vocabulary *is* applied to your transcription.

For examples of [LanguageIdSettings](#) in a request, refer to Option 2 in the **AWS CLI** and **AWS SDK** drop-down panels in the [Using language identification with batch transcriptions](#) (p. 162) section.

Using language identification with batch transcriptions

You can use automatic language identification in a batch transcription job using the **AWS Management Console**, **AWS CLI**, or **AWS SDK**; see the following for examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This opens the **Specify job details** page.
3. In the **Job settings** panel, find the **Language settings** section and select **Automatic language identification** or **Automatic multiple languages identification**.

You have the option to select multiple language options (from the *Select languages* drop-down box) if you know which languages are present in your audio file. Providing language options can improve accuracy, but is not required.

Specify job details Info

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Language settings

You can transcribe your audio file in a language that you specify or have Amazon Transcribe identify and transcribe it in the predominant language.

Specific language Info

If you know the language spoken in your source audio, choose this option to get the most accurate results. The options available for additional processing vary between languages.

Automatic language identification Info

If you don't know the language spoken in your audio files, choose this option. You have access to fewer options for additional processing than if you choose **Specific language**.

Automatic multiple languages identification Info

If there are multiple languages spoken in your audio files and you're not sure what these languages are, choose this option. This selection provides limited additional processing options compared to **Specific language**.

Language options for automatic language identification - optional

To improve accuracy, choose at least two languages spoken the most often in your audio library. Amazon Transcribe chooses from one of the languages you've specified to transcribe each audio file. Leave this field empty if you're unsure about which languages to select.

Select languages

- English, US (en-US)
- English, AU (en-AU)
- English, UK (en-GB)
- Hindi, IN (hi-IN)
- Spanish, US (es-US)

4. Fill in any other fields you wish to include on the **Specify job details** page, then select **Next**. This takes you to the **Configure job - optional** page.

Configure job - optional [Info](#)

Audio settings

- Audio identification** [Info](#)
Choose to split multi-channel audio into separate channels for transcription, or identify speakers in the input audio.
- Alternative results** [Info](#)
Enable to view more transcription results

Content removal
Content removal conceals information in the resulting transcript from your source audio file. Amazon Transcribe changes items in the transcript and does not modify the source audio.

- PII redaction** [Info](#)
Label the type of PII and also mask the content with the PII entity type in the transcription output. For example, (123) 456-7890 will be masked as [PHONE].
- Vocabulary filtering** [Info](#)
Vocabulary filtering can remove, mask or tag specified words in the final transcript.

Customization

- Custom vocabulary** [Info](#)
A custom vocabulary improves the accuracy of recognizing words and phrases specific to your use case.

Cancel Previous **Create job**

5. Select **Create job** to run your transcription job.

AWS CLI

This example uses the `start-transcription-job` command and `IdentifyLanguage` parameter. For more information, see [StartTranscriptionJob](#) and [LanguageIdSettings](#).

Option 1: Without the `language-id-settings` parameter. Use this option if you are **not** including a custom language model, custom vocabulary, or custom vocabulary filter in your request. `language-options` is optional, but recommended.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--identify-language=true \ (or --identify-multiple-languages=true) \  
--language-options "en-US" "hi-IN"
```

Option 2: With the `language-id-settings` parameter. Use this option if you are including a custom language model, custom vocabulary, or custom vocabulary filter in your request.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--identify-language=true \ (or --identify-multiple-languages=true)  
--language-options "en-US" "hi-IN" \  
--language-id-settings en-US=VocabularyName=my-en-US-vocabulary,en-  
US=VocabularyFilterName=my-en-US-vocabulary-filter,en-US=LanguageModelName=my-en-US-  
language-model,hi-IN=VocabularyName=my-hi-IN-vocabulary,hi-IN=VocabularyFilterName=my-hi-  
IN-vocabulary-filter
```

Here's another example using the `start-transcription-job` command, and a request body that identifies the language.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://filepath/my-first-language-id-job.json
```

The file `my-first-language-id-job.json` contains the following request body.

Option 1: Without the `LanguageIdSettings` parameter. Use this option if you are **not** including a custom language model, custom vocabulary, or custom vocabulary filter in your request. `LanguageOptions` is optional, but recommended.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "IdentifyLanguage": "true", (or "IdentifyMultipleLanguages": "true"),  
  "LanguageOptions": [  
    "en-US", "hi-IN"  
  ]  
}
```

Option 2: With the `LanguageIdSettings` parameter. Use this option if you are including a custom language model, custom vocabulary, or custom vocabulary filter in your request.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "IdentifyLanguage": "true", (or "IdentifyMultipleLanguages": "true")  
  "LanguageOptions": [  
    "en-US", "hi-IN"  
  ],  
  "LanguageIdSettings": {  
    "en-US": {  
      "LanguageModelName": "my-en-US-language-model",  
      "VocabularyFilterName": "my-en-US-vocabulary-filter",
```

```
    "VocabularyName": "my-en-US-vocabulary"  
  },  
  "hi-IN": {  
    "VocabularyName": "my-hi-IN-vocabulary",  
    "VocabularyFilterName": "my-hi-IN-vocabulary-filter"  
  }  
}  
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to identify your file's language using the `IdentifyLanguage` argument for the `start_transcription_job` method. For more information, see [StartTranscriptionJob](#) and [LanguageIdSettings](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

Option 1: Without the `LanguageIdSettings` parameter. Use this option if you are **not** including a custom language model, custom vocabulary, or custom vocabulary filter in your request. `LanguageOptions` is optional, but recommended.

```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe', 'us-west-2')  
job_name = "my-first-transcription-job"  
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
transcribe.start_transcription_job(  
    TranscriptionJobName = job_name,  
    Media = {  
        'MediaFileUri': job_uri  
    },  
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',  
    OutputKey = 'my-output-files/',  
    MediaFormat = 'flac',  
    IdentifyLanguage = True, (or IdentifyMultipleLanguages = True),  
    LanguageOptions = [  
        'en-US', 'hi-IN'  
    ]  
)  
  
while True:  
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)  
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:  
        break  
    print("Not ready yet...")  
    time.sleep(5)  
print(status)
```

Option 2: With the `LanguageIdSettings` parameter. Use this option if you are including a custom language model, custom vocabulary, or custom vocabulary filter in your request.

```
from __future__ import print_function  
import time  
import boto3  
transcribe = boto3.client('transcribe')  
job_name = "my-first-transcription-job"  
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
transcribe.start_transcription_job(  
    TranscriptionJobName = job_name,
```

```
Media = {
    'MediaFileUri': job_uri
},
OutputBucketName = 'DOC-EXAMPLE-BUCKET',
OutputKey = 'my-output-files/',
MediaFormat='flac',
IdentifyLanguage=True, (or IdentifyMultipleLanguages=True)
LanguageOptions = [
    'en-US', 'hi-IN'
],
LanguageIdSettings={
    'en-US': {
        'VocabularyName': 'my-en-US-vocabulary',
        'VocabularyFilterName': 'my-en-US-vocabulary-filter',
        'LanguageModelName': 'my-en-US-language-model'
    },
    'hi-IN': {
        'VocabularyName': 'my-hi-IN-vocabulary',
        'VocabularyFilterName': 'my-hi-IN-vocabulary-filter'
    }
}
)

while True:
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Language identification with streaming transcriptions

Streaming language identification can identify the dominant language spoken in your media stream. Amazon Transcribe requires a minimum of three seconds of speech to identify the language.

To use streaming language identification, you must provide at least two language codes, and you can select only one language dialect per language per stream. This means that you cannot select en-US and en-AU as language options for the same transcription.

You also have the option to select a preferred language from the set of language codes you provide. Adding a preferred language can speed up the language identification process, which is helpful for short audio clips.

Important

If none of the language codes you provide match the language, or languages, identified in your audio, Amazon Transcribe selects the closest language match from your specified language codes. It then produces a transcript in that language. For example, if your media is in US English (en-US) and you provide Amazon Transcribe with the language codes zh-CN, fr-FR, and de-DE, Amazon Transcribe is likely to match your media to German (de-DE) and produce a German-language transcription. Mismatching language codes and spoken languages can result in an inaccurate transcript, so we advise caution when including language codes.

If your media contains two channels, Amazon Transcribe can identify the dominant language spoken in each channel. In this case, set the [ChannelIdentification](#) parameter to `true` and each channel is transcribed separately. Note that the default for this parameter is `false`. If you don't change it, only the first channel is transcribed and only one language is identified.

Streaming language identification can't be combined with custom language models or redaction. If combining language identification with other features, you are limited to the languages supported with those features, and also with streaming transcriptions. Refer to [Supported languages and language-specific features](#) (p. 3) for more information.

Note

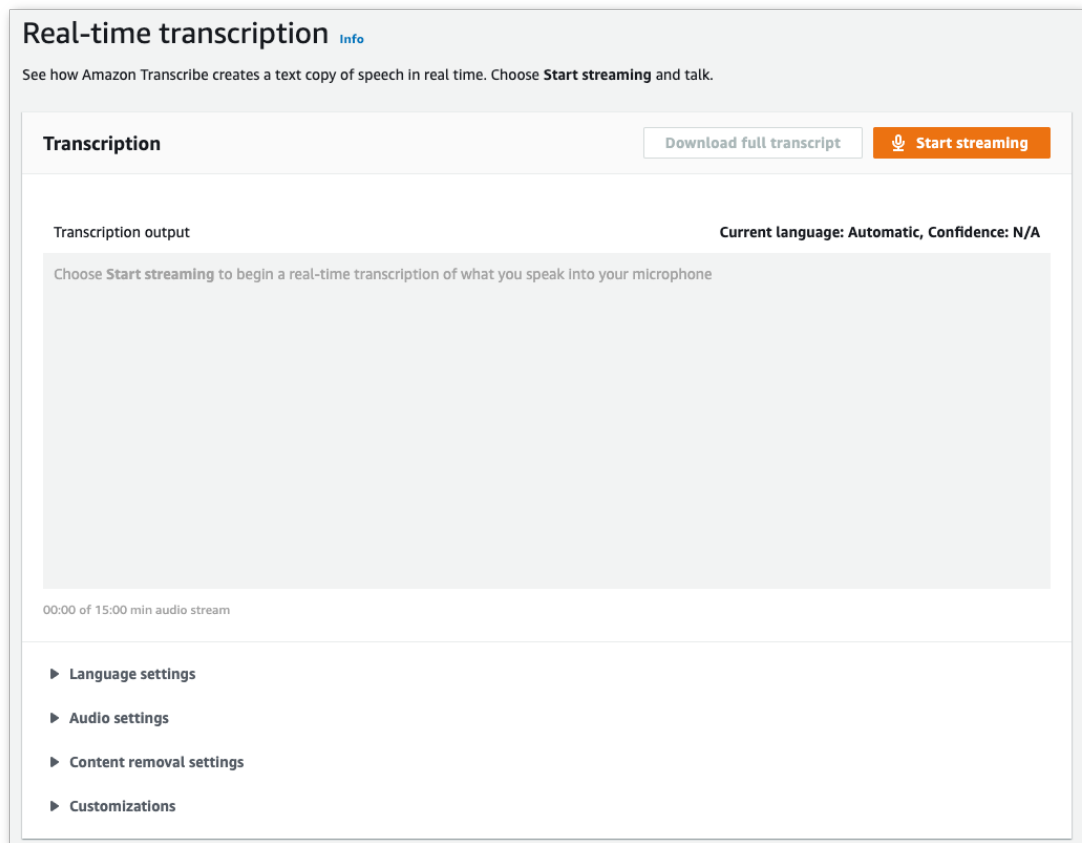
PCM and FLAC are the only supported audio formats for streaming language identification.

Using language identification with streaming media

You can use automatic language identification in a streaming transcription using the **AWS Management Console**, **HTTP/2**, or **WebSockets**; see the following for examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Real-time transcription**. Scroll down to **Language settings** and expand this field if it is minimized.



3. Select **Automatic language identification**.

▼ **Language settings**

Language settings
You can select a specific language for your transcription or have Amazon Transcribe identify the predominant language in your media and perform the transcription in that language.

Specific language
If you know the language spoken in your source audio, choose this option to get the most accurate results.

Automatic language identification [Info](#)
If you don't know the language spoken in your audio files, choose this option.

Language options for automatic language identification
To improve language identification accuracy, select a minimum of 2 language options.

Choose language(s) ▼

Preferred language - optional
Specify one preferred language from your previous selection.

Choose language ▼

▶ **Audio settings**

▶ **Content removal settings**

▶ **Customizations**

4. Provide a minimum of two language codes for your transcription. Note that you can provide only one dialect per language. For example, you cannot select both `en-US` and `en-AU` as language options for the same transcription.

Q

English, US (en-US)

Chinese, CN (zh-CN)

English, AU (en-AU)

English, UK (en-GB)

French, CA (fr-CA)

French, FR (fr-FR)

German, DE (de-DE)

Italian, IT (it-IT)

Japanese, JP (ja-JP)

Korean, KR (ko-KR)

Portuguese, BR (pt-BR)

Spanish, US (es-US)

Choose language(s) ▲

English, US (en-US) X French, CA (fr-CA) X

Preferred language - optional
Specify preferred language from the language options you selected

▼

5. (Optional) From the subset of languages you selected in the previous step, you can choose a preferred language for your transcript.

▼ Language settings

Language settings
You can select a specific language for your transcription or have Amazon Transcribe identify the predominant language in your media and perform the transcription in that language.

Specific language
If you know the language spoken in your source audio, choose this option to get the most accurate results.

Automatic language identification [Info](#)
If you don't know the language spoken in your audio files, choose this option.

Language options for automatic language identification
To improve language identification accuracy, select a minimum of 2 language options.

Choose language(s) ▼

English, US (en-US) X French, CA (fr-CA) X

Preferred language - optional
Specify one preferred language from your previous selection.

Choose language ▲

None

English, US (en-US)

French, CA (fr-CA)

► Customizations

6. You're now ready to transcribe your stream. Select **Start streaming** and begin speaking. To end your dictation, select **Stop streaming**.

HTTP/2 stream

This example creates an HTTP/2 request with language identification enabled. For more information on using HTTP/2 streaming with Amazon Transcribe, see [Setting up an HTTP/2 stream \(p. 82\)](#). For more detail on parameters and headers specific to Amazon Transcribe, see [StartStreamTranscription](#).

```
POST /stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
X-Amz-Target: com.amazonaws.transcribe.Transcribe.StartStreamTranscription
Content-Type: application/vnd.amazon.eventstream
X-Amz-Content-Sha256: string
X-Amz-Date: 20220208T235959Z
Authorization: AWS4-HMAC-SHA256 Credential=access-key/20220208/us-west-2/transcribe/
aws4_request, SignedHeaders=content-type;host;x-amz-content-sha256;x-amz-date;x-amz-
target;x-amz-security-token, Signature=string
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
x-amzn-transcribe-identify-language: true
x-amzn-transcribe-language-options: en-US,de-DE
x-amzn-transcribe-preferred-language: en-US
transfer-encoding: chunked
```

If you use `identify-language` in your request, you must also include `language-options`. You cannot use both `language-code` and `identify-language` in the same request.

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

WebSocket stream

This example creates a pre-signed URL that uses language identification in a WebSocket stream. Line breaks have been added for readability. For more information on using WebSocket streams with

Amazon Transcribe, see [Setting up a WebSocket stream \(p. 86\)](#). For more detail on parameters, see [StartStreamTranscription](#).

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/stream-transcription-websocket?
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=string
&X-Amz-SignedHeaders=content-type%3Bhost%3Bx-amz-date
&media-encoding=flac
&sample-rate=16000
&identify-language=true
&language-options=en-US,de-DE
&preferred-language=en-US
```

If you use `identify-language` in your request, you must also include `language-options`. You cannot use both `language-code` and `identify-language` in the same request.

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

Identifying speakers (diarization)

To identify different speakers in Amazon Transcribe, use *speaker diarization*. When you enable speaker diarization, Amazon Transcribe labels each speaker utterance. You enable speaker diarization by using the batch transcription or real-time streaming APIs, or the AWS Management Console.

Identifying speakers in audio

You can enable speaker diarization in a batch transcription job using either the [StartTranscriptionJob](#) API or the [AWS Management Console](#).

AWS Management Console

To use the AWS Management Console to enable speaker diarization in your transcription job, you enable audio identification and then speaker diarization.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, provide information about your transcription job.
5. Choose **Next**.
6. Enable **Audio identification**.
7. For **Audio identification type**, choose **Speaker identification**.
8. For **Maximum number of speakers**, specify the maximum number of speakers you think are speaking in your audio. For best results, match the number of speakers you ask Amazon Transcribe to identify to the number of speakers in the input audio. If you specify a value less than the number of speakers in your input audio, the transcription text of the most similar sounding speakers are attributed to a speaker label.
9. Choose **Create**.

AWS CLI

- Run the following code.

```
aws transcribe start-transcription-job \  
--cli-input-json file://filepath/example-start-command.json
```

The following code shows the contents of `example-start-command.json`.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "LanguageCode": "en-US",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
}
```

```
"OutputKey": "my-output-files/",  
"Settings":{  
  "MaxSpeakerLabels": 2,  
  "ShowSpeakerLabels":true  
}  
}
```

Identifying speakers in real-time streams

You can identify different speakers in either HTTP/2 or WebSocket streams. Speaker diarization works best for identifying between two and five speakers. Although Amazon Transcribe can identify more than five speakers in a stream, the accuracy of speaker diarization decreases if you exceed that number. To start an HTTP/2 stream, you specify the `ShowSpeakerLabel` request parameter of the [StartStreamTranscription](#) API. To start a WebSocket request, you use a pre-signed URI, which is a URI that contains the information needed to start your stream. To use the AWS Management Console to transcribe speech spoken into your microphone, use the following procedure.

You can identify speakers in real-time streams that are in US English (en-US).

You can use the [AWS Management Console](#) to start a real-time stream and transcribe any speech picked up by your microphone.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Real-time transcription**.
3. In **Language**, choose the language of your real-time stream.
4. Under **Additional settings**, enable **Speaker identification**.
5. Choose **Start streaming**.
6. Speak into the microphone.

HTTP/2 streaming

The following is the syntax for the parameters of an HTTP/2 request.

```
POST /stream-transcription HTTP/2  
host: transcribestreaming.us-west-2.amazonaws.com  
authorization: Generated value  
x-amz-target: com.amazonaws.transcribe.Transcribe.StartStreamTranscription  
x-amz-content-sha256: STREAMING-AWS4-HMAC-SHA256-EVENTS  
x-amz-date: 20220208T235959Z  
x-amzn-transcribe-session-id: my-first-http2-diarization-stream  
x-amzn-transcribe-language-code: en-US  
x-amzn-transcribe-media-encoding: flac  
x-amzn-transcribe-sample-rate: 16000  
x-amzn-transcribe-show-speaker-label: true  
content-type: application/vnd.amazon.eventstream  
transfer-encoding: chunked
```

Parameter descriptions:

- **host:** Update the AWS Region ('us-west-2' in the preceding example) with the AWS Region you are calling. For a list of valid AWS Regions, see [AWS Regions and Endpoints](#).
- **authorization:** This is a generated field. To learn more about creating a signature, see [Signing AWS requests with Signature Version 4](#).
- **x-amz-target:** Don't alter this field; use the content shown in the preceding example.

- **x-amz-content-sha256**: This is a generated field. To learn more about calculating a signature, see [Signing AWS requests with Signature Version 4](#).
- **x-amz-date**: The date and time the signature is created. The format is YYYYMMDDTHHMMSSZ, where YYYY=year, MM=month, DD=day, HH=hour, MM=minute, SS=seconds, and 'T' and 'Z' are fixed characters. For more information, refer to [Handling Dates in Signature Version 4](#).
- **x-amzn-transcribe-session-id**: The name for your streaming session.
- **x-amzn-transcribe-language-code**: The encoding used for your input audio. Refer to [StartStreamTranscription](#) or [Supported languages and language-specific features \(p. 3\)](#) for a list of valid values.
- **x-amzn-transcribe-media-encoding**: The encoding used for your input audio. Valid values are pcm, ogg-opus, and flac.
- **x-amzn-transcribe-sample-rate**: The sample rate of the input audio (in Hertz). Amazon Transcribe supports a range from 8,000 Hz to 48,000 Hz. Low-quality audio, such as telephone audio, is typically around 8,000 Hz. High-quality audio typically ranges from 16,000 Hz to 48,000 Hz. Note that the sample rate you specify **must** match that of your audio.
- **x-amzn-transcribe-show-speaker-label**: To enable diarization, this value must be true.
- **content-type**: Don't alter this field; use the content shown in the preceding example.

WebSocket streaming

To identify speakers in WebSocket streams, use the following format to create a pre-signed URI to start a WebSocket request and specify `show-speaker-label` as true. A pre-signed URI contains the information to set up bi-directional communication between your application and Amazon Transcribe.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/stream-transcription-websocket
?
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=Signature Version 4 signature
&X-Amz-SignedHeaders=host
&language-code=en-US
&media-encoding=flac
&sample-rate=16000
&show-speaker-label=true
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

Creating video subtitles

Amazon Transcribe supports WebVTT (*.vtt) and SubRip (*.srt) output for use as video subtitles. You can select one or both file types when setting up your batch video transcription job. When using the subtitle feature, your selected subtitle file(s) and a regular transcript file (containing additional information) are produced. Subtitle and transcription files are output to the same destination.

Subtitles are displayed at the same time text is spoken, and remain visible until there is a natural pause or the speaker finishes talking.

Important

Amazon Transcribe uses a default start index of 0 for subtitle output, which differs from the more widely used value of 1. If you require a start index of 1, you can specify this in the AWS Management Console or in your API request using the [OutputStartIndex](#) parameter.

Using the incorrect start index can result in compatibility errors with other services, so be sure to verify which start index you require before creating your subtitles. If you're uncertain which value to use, we recommend choosing 1. Refer to [Subtitles](#) for more information.

Features supported with subtitles:

- **Content redaction** — Any redacted content is reflected as 'PII' in both your subtitle and regular transcript output files. The audio is not altered.
- **Vocabulary filters** — Subtitle files are generated from the transcription file, so any words you filter in your standard transcription output are also filtered in your subtitles. Filtered content is displayed as whitespace or *** in your transcript and subtitle files. The audio is not altered.
- **Speaker diarization** — If there are multiple speakers in a given subtitle segment, dashes are used to distinguish each speaker. This applies to both WebVTT and SubRip formats; for example:
 - -- Text spoken by Person 1
 - -- Text spoken by Person 2

Subtitle files are stored in the same Amazon S3 location as your transcription output.

See [Amazon Transcribe Video Snacks: Creating Video Subtitles Without Writing Any Code](#) for a video walkthrough of creating subtitles.

Generating subtitle files

You can create subtitle files using the [AWS Management Console](#), [AWS SDK](#), or [AWS CLI](#); see the following examples:

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Transcription jobs**, then select **Create job** (top right). This opens the **Specify job details** page. Subtitle options are located in the **Output data** panel.
3. Select the formats you want for your subtitle files, then choose a value for your start index. Note that the Amazon Transcribe default is 0, but 1 is more widely used. If you're uncertain which value to use, we recommend choosing 1, as this may improve compatibility with other services.

Output data

Output data location type info [Info](#)

Service-managed S3 bucket
The output will be removed after 90 days when the job expires.

Customer specified S3 bucket
The output will not be removed from bucket even after the job expires.

Subtitle file format [Info](#)

SRT (SubRip)

VTT (WebVTT)

Specify the start index

0 ▼

4. Fill in any other fields you wish to include on the **Specify job details** page, then select **Next**. This takes you to the **Configure job - optional page**.
5. Select **Create job** to run your transcription job.

AWS CLI

This example uses the `start-transcription-job` command and `Subtitles` parameter. For more information, see [StartTranscriptionJob](#) and [Subtitles](#).

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--transcription-job-name my-first-transcription-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--output-key my-output-files/ \  
--language-code en-US \  
--subtitles Formats=vtt,srt,OutputStartIndex=1
```

Here's another example using the `start-transcription-job` command, and a request body that adds subtitles to that job.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://my-first-subtitle-job.json
```

The file `my-first-subtitle-job.json` contains the following request body.

```
{  
  "TranscriptionJobName": "my-first-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "LanguageCode": "en-US",  
  "Subtitles": {  
    "Formats": [  
      "vtt", "srt"  
    ],  
  },  
}
```

```
        "OutputStartIndex": 1
    }
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to add subtitles using the `Subtitles` argument for the `start_transcription_job` method. For more information, see [StartTranscriptionJob](#) and [Subtitles](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
transcribe.start_transcription_job(
    TranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'my-output-files/',
    LanguageCode = 'en-US',
    Subtitles = {
        'Formats': [
            'vtt', 'srt'
        ],
        'OutputStartIndex': 1
    }
)

while True:
    status = transcribe.get_transcription_job(TranscriptionJobName = job_name)
    if status['TranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Analyzing call center audio with Call Analytics

Use Amazon Transcribe Call Analytics to gain insight into customer-agent interactions. Call Analytics provides you with:

- [Call characteristics \(p. 180\)](#), including talk time, non-talk time, speaker loudness, interruptions, and talk speed
- [Call summarization \(p. 180\)](#), which detects issues, action items, and outcomes
- [Custom categories \(p. 179\)](#) and rules that you can use to hone in on specific keywords and criteria
- [Redaction \(p. 181\)](#) of your text transcript and your audio file
- [Speaker sentiment \(p. 181\)](#) for each caller at various points in a call

Call Analytics can also help increase staff productivity by reducing the need for manual note-taking after each call. It also reduces the time required by supervisors to read transcripts or listen to audio recordings.

Call Analytics use cases

- **Monitor issue frequency over time:** Use [call categorization \(p. 179\)](#) to identify recurring keywords within your transcripts.
- **Gain insight into your customer service experience:** Use [call characteristics \(p. 180\)](#) (non-talk time, talk time, interruptions, voice loudness, talk speed) and sentiment analysis to determine if customers' issues are being appropriately resolved during their call.
- **Ensure regulatory compliance or adherence to company policy:** Set [keywords and phrases \(p. 179\)](#) for company-specific greetings or disclaimers to verify your agents are meeting regulatory requirements.
- **Improve handling of customer data:** Use transcript and audio [redaction \(p. 181\)](#) to protect customer privacy.
- **Improve staff training:** Use criteria (sentiment, non-talk time, interruptions, talk speed) to flag transcripts that can be used as examples of positive or negative customer interactions.
- **Measure staff efficacy in creating a positive customer experience:** Use [sentiment analysis \(p. 181\)](#) to measure if your agents are able to turn a negative customer sentiment into a positive one as calls progress.
- **Improve data organization:** Label and sort calls based on [custom categories \(p. 179\)](#) (including keywords and phrases, sentiment, talk time, and interruptions).
- **Quickly summarize the important aspects of a call** without needing to review the entire transcript: Use automatic [call summarization \(p. 180\)](#) to get a succinct summary of all issues, action items, and outcomes identified in a given call.

To compare the features available with Call Analytics to those for Amazon Transcribe and Amazon Transcribe Medical, refer to the [Amazon Transcribe features \(p. 40\)](#) table.

To get started, see [Starting a Call Analytics job \(p. 189\)](#). Call Analytics output is similar to that of a standard transcription job, but contains additional analytics data. To view Call Analytics output, see [Example output \(p. 194\)](#).

API operations specific to Call Analytics

[CreateCallAnalyticsCategory](#), [DeleteCallAnalyticsCategory](#), [DeleteCallAnalyticsJob](#), [GetCallAnalyticsCategory](#), [GetCallAnalyticsJob](#), [ListCallAnalyticsCategories](#), [ListCallAnalyticsJobs](#), [StartCallAnalyticsJob](#), [UpdateCallAnalyticsCategory](#)

Considerations and additional information

Before using Call Analytics, note that:

- Call Analytics only supports stereo-channel audio; mono-channel audio cannot be used with this feature.
- [Job queueing \(p. 92\)](#) is always enabled, meaning that you are limited to 100 concurrent Call Analytics jobs. If you need to request a quota increase, see [AWS service quotas](#).
- Input files for Call Analytics cannot be greater than 500 MB and must be less than 4 hours.
- If using categories, you must create all desired categories before starting a Call Analytics job. Any new categories cannot be applied to existing Call Analytics jobs. To learn how to create a new category, see [Creating categories \(p. 181\)](#).
- If you're using redaction, only US English (en-US) is supported. Refer to [Supported languages and language-specific features \(p. 3\)](#) for a complete list of languages supported with Call Analytics.
- Call summarization is only supported for the following English language dialects: Australian (en-AU), British (en-GB), Indian (en-IN), Irish (en-IE), Scottish (en-AB), US (en-US), and Welsh (en-WL).
- Some Call Analytics quotas differ from Amazon Transcribe and Amazon Transcribe Medical; refer to [Guidelines and quotas \(p. 42\)](#) for specifics.

Tip

To learn more about post-Call Analytics options, see [Post Call Analytics for your contact center with Amazon language AI services](#).

To view sample Call Analytics output and features, check out our [GitHub demo](#). We also have a [JSON to Word document GitHub example](#) for Call Analytics.

Call Analytics insights

The insights available with Call Analytics are separated into five main groups:

- [Call categorization \(p. 179\)](#)
- [Call characteristics \(p. 180\)](#)
- [Call summarization \(p. 180\)](#)
- [Sensitive data redaction \(p. 181\)](#)
- [Sentiment analysis \(p. 181\)](#)

Call categorization

Use call categorization to flag keywords, phrases, sentiment, or actions during a call. Our categorization options can help you triage escalations, such as negative-sentiment calls with many interruptions, or organize calls into specific categories, such as company departments.

The criteria you can add to a category include:

- **Non-talk time:** Periods of time when neither the customer nor the agent is talking.
- **Interruptions:** When either the customer or the agent is interrupting the other person.
- **Customer or agent sentiment:** How either the customer or the agent is feeling during a specified time period. If at least 50 percent of the conversation turns (the back-and-forth between two speakers) in a specified time period match the specified sentiment, Amazon Transcribe considers the sentiment a match.
- **Keywords or phrases:** Matches part of the transcription based on an exact phrase. For example, if you set a filter for the phrase "I want to speak to the manager", Amazon Transcribe filters for that *exact* phrase.

Here's an [output example \(p. 194\)](#).

For more information on categories or to learn how to create a new category, see [Creating categories \(p. 181\)](#).

Call characteristics

The call characteristics feature measures the quality of agent-customer interactions using these criteria:

- **Interruption**—Measures if and when one participant cuts off the other participant mid-sentence. Frequent interruptions may be associated with rudeness or anger, and could correlate to negative sentiment for one or both participants.
- **Loudness**—Measures the volume at which each participant is speaking. Use this metric to see if either the caller or the agent is speaking loudly or yelling, which is often indicative of being upset. This metric is represented as a normalized value (speech level per second of speech in a given segment) on a scale from 0 to 100, where a higher value indicates a louder voice.
- **Non-talk time**—Measures periods of time that do not contain speech. Use this metric to see if there are long periods of silence, such as an agent keeping a customer on hold for an excessive amount of time.
- **Talk speed**—Measures the speed at which both participants are speaking. Comprehension can be affected if one participant speaks too quickly. This metric is measured in words per minute.
- **Talk time**—Measures the amount of time (in milliseconds) each participant spoke during the call. Use this metric to help identify if one participant is dominating the call or if the dialogue is balanced.

Here's an [output example \(p. 194\)](#).

Call summarization

Call summarization provides succinct summaries of the important components in agent-customer calls, including issues, action items, and outcomes for each participant.

Using call summarization, you can:

- Reduce the need for manual note-taking during and after calls
- Improve agent efficiency, allowing them to respond faster to customers
- Increase customer satisfaction through a reduction in hold times
- Streamline supervisor reviews, as call summaries are much quicker to digest than entire transcripts

Call summarization works across all industries and business sectors, and is context-based.

Note

Call summarization is supported with these English language dialects: Australian (en-AU), British (en-GB), Indian (en-IN), Irish (en-IE), Scottish (en-AB), US (en-US), and Welsh (en-WL).

Call summarization works out-of-the-box, and thus doesn't support customization, such as model training or custom categories.

Here's an [output example \(p. 195\)](#).

Sensitive data redaction

Sensitive data redaction replaces personally identifiable information (PII) in the text transcript and the audio file. A redacted transcript replaces the original text with [P I I]; a redacted audio file replaces spoken personal information with silence. This parameter is useful for protecting customer information.

Note

Redaction is supported with US English (en-US) only.

PII redaction works out-of-the-box, and thus doesn't support customization. To view the list of PII that is redacted using this feature, or to learn more about redaction with Amazon Transcribe, see [Redacting or identifying personally identifiable information \(p. 142\)](#).

Here's an [output example \(p. 196\)](#).

Sentiment analysis

Sentiment analysis estimates how the customer and agent are feeling throughout the call. This metric is represented as both a quantitative value (with a range from 5 to -5) and a qualitative value (positive, neutral, mixed, or negative). Quantitative values are provided per quarter and per call; qualitative values are provided per turn.

Using this parameter, you can quantitatively evaluate the overall sentiment for each call participant and the sentiment for each participant during each quarter of the call. You can qualitatively evaluate turn-by-turn sentiment for each participant. This metric can help identify if your agent is able to delight an upset customer by the time the call ends.

Sentiment analysis works out-of-the-box, and thus doesn't support customization, such as model training or custom categories.

Here's an [output example \(p. 196\)](#).

Creating categories

Call Analytics supports the creation of custom categories so you can tailor your transcript analysis to your specific business needs.

You can create as many categories as you'd like to cover a range of different scenarios. For each category you create, you must create between 1 and 20 rules. Each rule is based on one of four criteria: interruptions, keywords, non-talk time, or sentiment. For more details on using these criteria with the [CreateCallAnalyticsCategory](#) operation, refer to the [Rule criteria \(p. 187\)](#) section.

Here are a few examples of what you can do with custom categories:

- Isolate calls with specific characteristics, such as calls that end with a negative customer sentiment


- Identify trends in customer issues by flagging and tracking specific sets of keywords
- Monitor compliance, such as an agent speaking a specific phrase during the first few seconds of a call
- Gain insight into customer experience by flagging calls with many agent interruptions and negative customer sentiment
- Compare multiple categories to measure correlations, such as analyzing whether an agent using a welcome phrase correlates with positive customer sentiment

If you run a Call Analytics job and the call characteristics match the rules you've specified in a category, Amazon Transcribe labels your job with that category. See [call categorization output \(p. 194\)](#) for an example of a category match in JSON output.

To create a category, you can use the **AWS Management Console**, **AWS CLI**, or **AWS SDK**; see the following for examples:

AWS Management Console

1. In the navigation pane, under Amazon Transcribe, choose **Amazon Transcribe Call Analytics**.
2. Choose **Call analytics categories**, which takes you to the **Call analytics categories** page. Select **Create category**.



3. You're now on the **Create category** page. Enter a name for your category, then choose if you're going to use a template to create a category or make one from scratch.

If using a template: select **Use a template (recommended)**, choose the template you want, then select **Create category**.

4. If creating a custom category: select **Create from scratch**.

Create category [Info](#)

Category settings

Category name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, -, _ , and - (hyphen).

Category creation method [Info](#)

Use a template (recommended)
Use a template to edit predefined rules.

Create from scratch
If you know the rules that you want to define, choose this option.

Rules
All the rule conditions must be met for a transcription job to be classified in this category.

▼ Rule 1 Delete rule

Rule type [Info](#)
Choose the rule that you want to define.

Add rule
You can add up to 19 more rules.

5. Add rules to your category using the drop-down menu. You can add up to 20 rules per category.

Rules
All the rule conditions must be met for a transcription job to be classified in this category.

▼ Rule 1 Delete rule

When no word has been spoken for more than 5 minute(s) during the entire call.

Rule type [Info](#)
Choose the rule that you want to define.

Non-talk time ▲
Trigger the rule when nothing has been said for more than the time that you specify.

Interruption time
Trigger the rule when the speaker has been interrupted for more than the time that you specify.

Transcript content match
Trigger the rule when the speaker says the words or phrases that you specify.

Transcript sentiment match
Trigger the rule when the sentiment of the speaker matches the sentiment that you specify.

Add rule
You can add up to 19 more rules.

6. Here's an example of a category with two rules: an agent who interrupts a customer for more than 15 seconds during the call and there is negative sentiment by either the customer or the agent in the last two minutes of the call.

Rules
All the rule conditions must be met for a transcription job to be classified in this category.

▼ **Rule 1** Delete rule

When the duration of the interruption was more than 15 second(s) during the entire call when the speaker was agent.

Rule type [Info](#)
Choose the rule that you want to define.
Interruption time ▼

Logic [Info](#)
Define the conditions that must be met.

When the duration of the interruption was more than 15 second(s) ▼
during the entire call ▼
when the speaker was agent ▼

AND

▼ **Rule 2** Delete rule

When the sentiment is negative during the last 2 minute(s) when the speaker was either.

Rule type [Info](#)
Choose the rule that you want to define.
Transcript sentiment match ▼

Logic [Info](#)
Define the conditions that must be met.

When the sentiment is negative ▼
during the last ▼ 2 minute(s) ▼
when the speaker was either ▼

Add rule
You can add up to 18 more rules.

7. When you're finished adding rules to your category, choose **Create category**.

AWS CLI

This example uses the `create-call-analytics-category` command. For more information, see [CreateCallAnalyticsCategory](#), [CategoryProperties](#), and [Rule](#).

The following example creates a category with the rules:

- The customer was interrupted in the first 60,000 milliseconds. The duration of these interruptions lasted at least 10,000 milliseconds.
- There was a period of silence that lasted at least 20,000 milliseconds between 10% into the call and 80% into the call.
- The agent had a negative sentiment at some point in the call.
- The words "welcome" or "hello" were not used in the first 10,000 milliseconds of the call.

This example uses the `create-call-analytics-category` command, and a request body that adds several rules to your category.

```
aws transcribe create-call-analytics-category \  
--cli-input-json file://filepath/my-first-analytics-category.json
```

The file `my-first-analytics-category.json` contains the following request body.

```
{  
  "CategoryName": "my-new-category",  
  "Rules": [  
    {  
      "InterruptionFilter": {  
        "AbsoluteTimeRange": {  
          "First": 60000  
        },  
        "Negate": false,  
        "ParticipantRole": "CUSTOMER",  
        "Threshold": 10000  
      },  
    },  
    {  
      "NonTalkTimeFilter": {  
        "Negate": false,  
        "RelativeTimeRange": {  
          "EndPercentage": 80,  
          "StartPercentage": 10  
        },  
        "Threshold": 20000  
      },  
    },  
    {  
      "SentimentFilter": {  
        "ParticipantRole": "AGENT",  
        "Sentiments": [  
          "NEGATIVE"  
        ]  
      },  
    },  
    {  
      "TranscriptFilter": {  
        "Negate": true,  
        "AbsoluteTimeRange": {  
          "First": 10000  
        },  
        "Targets": [  
          "welcome",  
          "hello"  
        ],  
        "TranscriptFilterType": "EXACT"  
      },  
    }  
  ]  
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to create a category using the `CategoryName` and `Rules` arguments for the `create_call_analytics_category` method. For more information, see [CreateCallAnalyticsCategory](#), [CategoryProperties](#), and [Rule](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

The following example creates a category with the rules:

- The customer was interrupted in the first 60,000 milliseconds. The duration of these interruptions lasted at least 10,000 milliseconds.
- There was a period of silence that lasted at least 20,000 milliseconds between 10% into the call and 80% into the call.
- The agent had a negative sentiment at some point in the call.
- The words "welcome" or "hello" were not used in the first 10,000 milliseconds of the call.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
category_name = "my-new-category"
transcribe.create_call_analytics_category(
    CategoryName = category_name,
    Rules = [
        {
            'InterruptionFilter': {
                'AbsoluteTimeRange': {
                    'First': 60000
                },
                'Negate': False,
                'ParticipantRole': 'CUSTOMER',
                'Threshold': 10000
            }
        },
        {
            'NonTalkTimeFilter': {
                'Negate': False,
                'RelativeTimeRange': {
                    'EndPercentage': 80,
                    'StartPercentage': 10
                },
                'Threshold': 20000
            }
        },
        {
            'SentimentFilter': {
                'ParticipantRole': 'AGENT',
                'Sentiments': [
                    'NEGATIVE'
                ]
            }
        },
        {
            'TranscriptFilter': {
                'Negate': True,
                'AbsoluteTimeRange': {
                    'First': 10000
                },
                'Targets': [
                    'welcome',
                    'hello'
                ],
                'TranscriptFilterType': 'EXACT'
            }
        }
    ]
)

result = transcribe.get_call_analytics_category(CategoryName = category_name)
```



```
print(result)
```

Rule criteria

This section outlines the types of custom rules you can create using the [CreateCallAnalyticsCategory](#) API operation.

Interruption match

Rules using interruptions ([InterruptionFilter](#) data type) are designed to match:

- Instances where an agent interrupts a customer
- Instances where a customer interrupts an agent
- Any participant interrupting the other
- A lack of interruptions

Here's an example of the parameters available with [InterruptionFilter](#):

```
"InterruptionFilter": {
  "AbsoluteTimeRange": {
    Specify the time frame, in milliseconds, when the match should occur
  },
  "RelativeTimeRange": {
    Specify the time frame, in percentage, when the match should occur
  },
  "Negate": Specify if you want to match the presence or absence of interruptions,
  "ParticipantRole": Specify if you want to match speech from the agent, the customer, or both,
  "Threshold": Specify a threshold for the amount of time, in seconds, interruptions occurred during the call
},
```

Refer to [CreateCallAnalyticsCategory](#) and [InterruptionFilter](#) for more information on these parameters and the valid values associated with each.

Keyword match

Rules using keywords ([TranscriptFilter](#) data type) are designed to match:

- Custom words or phrases spoken by the agent, the customer, or both
- Custom words or phrases **not** spoken by the agent, the customer, or both
- Custom words or phrases that occur at a specific time frame

Here's an example of the parameters available with [TranscriptFilter](#):

```
"TranscriptFilter": {
  "AbsoluteTimeRange": {
    Specify the time frame, in milliseconds, when the match should occur
  },
  "RelativeTimeRange": {
    Specify the time frame, in percentage, when the match should occur
  },
  "Negate": Specify if you want to match the presence or absence of your custom keywords,
```

```
"ParticipantRole": Specify if you want to match speech from the agent, the customer, or both,
"Targets": [ The custom words and phrases you want to match ],
"TranscriptFilterType": Use this parameter to specify an exact match for the specified targets
}
```

Refer to [CreateCallAnalyticsCategory](#) and [TranscriptFilter](#) for more information on these parameters and the valid values associated with each.

Non-talk time match

Rules using non-talk time ([NonTalkTimeFilter](#) data type) are designed to match:

- The presence of silence at specified periods throughout the call
- The presence of speech at specified periods throughout the call

Here's an example of the parameters available with [NonTalkTimeFilter](#):

```
"NonTalkTimeFilter": {
  "AbsoluteTimeRange": {
    Specify the time frame, in milliseconds, when the match should occur
  },
  "RelativeTimeRange": {
    Specify the time frame, in percentage, when the match should occur
  },
  "Negate": Specify if you want to match the presence or absence of speech,
  "Threshold": Specify a threshold for the amount of time, in seconds, silence (or speech) occurred during the call
},
```

Refer to [CreateCallAnalyticsCategory](#) and [NonTalkTimeFilter](#) for more information on these parameters and the valid values associated with each.

Sentiment match

Rules using sentiment ([SentimentFilter](#) data type) are designed to match:

- The presence or absence of a positive sentiment felt by the customer, agent, or both at specified points in the call
- The presence or absence of a negative sentiment felt by the customer, agent, or both at specified points in the call
- The presence or absence of a neutral sentiment felt by the customer, agent, or both at specified points in the call
- The presence or absence of a mixed sentiment felt by the customer, agent, or both at specified points in the call

Here's an example of the parameters available with [SentimentFilter](#):

```
"SentimentFilter": {
  "AbsoluteTimeRange": {
    Specify the time frame, in milliseconds, when the match should occur
  },
  "RelativeTimeRange": {
    Specify the time frame, in percentage, when the match should occur
  }
}
```

```
},  
  "Negate": Specify if you want to match the presence or absence of your chosen  
sentiment,  
  "ParticipantRole": Specify if you want to match speech from the agent, the customer, or  
both,  
  "Sentiments": [ The sentiments you want to match ]  
},
```

Refer to [CreateCallAnalyticsCategory](#) and [SentimentFilter](#) for more information on these parameters and the valid values associated with each.

Starting a Call Analytics job

Before running a Call Analytics transcription job, you must create all the [categories \(p. 181\)](#) you want Amazon Transcribe to match in your audio.

Note

Call Analytics jobs can't be retroactively matched to new categories. Only the categories you create *before* running a Call Analytics job can be applied to that job.


If you've created one or more categories, and your audio matches all the rules within at least one of your categories, Amazon Transcribe flags your output with the matching category. If you choose not to use categories, or if your audio doesn't match the rules specified in your categories, your transcript isn't flagged.

To start a Call Analytics job, you can use the **AWS Management Console**, **AWS CLI**, or **AWS SDK**; see the following for examples:

AWS Management Console

Use the following procedure to start a Call Analytics job. The calls that match all characteristics defined by a category are labeled with that category.

1. In the navigation pane, under Amazon Transcribe Call Analytics, choose **Call analytics jobs**.
2. Choose **Create job**.



3. On the **Specify job details** page, provide information about your Call Analytics job, including the location of your input data.

Specify job details [Info](#)

Job settings

Name

The name can be up to 200 characters long. Valid characters are a-z, A-Z, 0-9, . (period), _ (underscore), and - (hyphen).

Model type [Info](#)
Choose the type of model to use for the transcription job.

General model
To use a model that is not specialized for a particular use case, choose this option. Configuration options vary between languages.

Custom language model
To use a model that you trained for your specific use case, choose this option. This model has fewer configuration options than the general model.

Language settings
You can transcribe your audio file in a language that you specify or have Amazon Transcribe identify and transcribe it in the predominant language.

Specific language [Info](#)
If you know the language spoken in your source audio, choose this option to get the most accurate results. The options available for additional processing vary between languages.

Automatic language identification [Info](#)
If you don't know the language spoken in your audio files, choose this option. You have access to fewer options for additional processing than if you choose **Specific language**.

Language
Choose the language of the input audio.

Input data [Info](#)

Input file location on S3
Choose an input audio or video file in Amazon S3.

Valid file formats: MP3, MP4, WAV, FLAC, AMR, OGG, and WebM.

Agent audio channel identification [Info](#)
Choose the channel that has the speech from the agent. The other channel is used for the customer's speech.

Specify the desired Amazon S3 location of your output data and which IAM role to use.

Output data

Output data location type info [Info](#)

Service-managed S3 bucket
The output will be removed after 90 days when the job expires.

Customer specified S3 bucket
The output will not be removed from bucket even after the job expires.

Access permissions

IAM role [Info](#)

Use an existing IAM role

Create an IAM role
By choosing **Create job** you are authorizing creation of this role.

Permissions to access
Your role has access to these resources. The KMS key permission is used only if your input bucket is encrypted

Input S3 bucket and KMS decrypt permission to input bucket

Any S3 bucket and any KMS keys

Role name
Roles are prefixed with "AmazonTranscribeServiceRoleFullAccess-". Your newly created role has full access to the S3 bucket and KMS key for your account.

The name can be up to 64 characters long

▼ Role permissions details

Your new role has these permissions to give Amazon Transcribe access to the resources that you've specified.

Service	Access level	Resource
S3	List, Read, Write	All resources
Key Management Service	GenerateDataKey, Decrypt	All resources

Cancel **Next**

4. Choose **Next**.
5. For **Configure job**, turn on any optional features you want to include with your Call Analytics job.

Configure job - *optional* [Info](#)

Content removal

Content removal conceals information in the resulting transcript from your source audio file. Amazon Transcribe changes items in the transcript and does not modify the source audio.

Automatic content redaction [Info](#)
Automatic content redaction removes personally identifiable information (PII) in your transcripts. Redactions in transcripts show up as [PII].

Vocabulary filtering [Info](#)
Vocabulary filtering can remove, mask or tag specified words in the final transcript.

Customization

Custom vocabulary [Info](#)
A custom vocabulary improves the accuracy of recognizing words and phrases specific to your use case.

Categories

Create categories to classify calls. For example, you can create a category for all cancellation requests. When you run an analytics job, Amazon Transcribe applies that category to all calls that request cancellation.

You haven't defined any categories. [Create a new category.](#) [↗](#)

[Cancel](#) [Previous](#) [Create job](#)

6. Choose **Create job**.

AWS CLI

This example uses the `start-call-analytics-job` command and `channel-definitions` parameter. For more information, see [StartCallAnalyticsJob](#) and [ChannelDefinition](#).

```
aws transcribe start-call-analytics-job \  
--region us-west-2 \  
--call-analytics-job-name my-first-call-analytics-job \  
--media MediaFileUri=s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac \  
--output-location s3://DOC-EXAMPLE-BUCKET/my-output-files/ \  
--data-access-role-arn arn:aws:iam::111122223333:role/ExampleRole \  
--channel-definitions ChannelId=0,ParticipantRole=AGENT \  
ChannelId=1,ParticipantRole=CUSTOMER
```

Here's another example using the `start-call-analytics-job` command, and a request body that identifies audio channels for that job.

```
aws transcribe start-call-analytics-job \  
--region us-west-2 \  
--cli-input-json file://filepath/my-call-analytics-job.json
```

The file `my-call-analytics-job.json` contains the following request body.

```
{
  "CallAnalyticsJobName": "my-first-call-analytics-job",
  "DataAccessRoleArn": "arn:aws:iam::111122223333:role/ExampleRole",
  "Media": {
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
  },
  "OutputLocation": "s3://DOC-EXAMPLE-BUCKET/my-output-files/",
  "ChannelDefinitions": [
    {
      "ChannelId": 0,
      "ParticipantRole": "AGENT"
    },
    {
      "ChannelId": 1,
      "ParticipantRole": "CUSTOMER"
    }
  ]
}
```

AWS SDK for Python (Boto3)

This example uses the AWS SDK for Python (Boto3) to start a Call Analytics job using the `start_call_analytics_job` method. For more information, see [StartCallAnalyticsJob](#) and [ChannelDefinition](#).

For additional examples using the AWS SDKs, including feature-specific, scenario, and cross-service examples, refer to the [Code examples for Amazon Transcribe using AWS SDKs \(p. 202\)](#) chapter.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-call-analytics-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
output_location = "s3://DOC-EXAMPLE-BUCKET/my-output-files/"
data_access_role = "arn:aws:iam::111122223333:role/ExampleRole"
transcribe.start_call_analytics_job(
    CallAnalyticsJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    DataAccessRoleArn = data_access_role,
    OutputLocation = output_location,
    ChannelDefinitions = [
        {
            'ChannelId': 0,
            'ParticipantRole': 'AGENT'
        },
        {
            'ChannelId': 1,
            'ParticipantRole': 'CUSTOMER'
        }
    ]
)

while True:
    status = transcribe.get_call_analytics_job(CallAnalyticsJobName = job_name)
    if status['CallAnalyticsJob']['CallAnalyticsJobStatus'] in ['COMPLETED', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Example output

Call Analytics transcripts are displayed in a turn-by-turn format. Each segment includes data specific to that segment, which can include detected issues, action items, detected outcomes, loudness scores, redaction, and sentiment. Additionally, a summary of call characteristics is provided at the end of the transcript.

To increase accuracy and further customize your transcripts to your use case, such as including industry-specific terms, use [custom vocabularies \(p. 101\)](#) or [custom language models \(p. 115\)](#) with your Call Analytics job. To mask, remove, or tag words you don't want in your transcription results, such as profanity, add [vocabulary filtering \(p. 131\)](#).

The following sections show examples of redacted JSON output, separated by Call Analytics insight. An abridged, but otherwise complete, transcript is provided following the insight-specific examples.

Call categorization

Here's what a category match looks like in your transcription output. This example shows that the audio from the 40040 millisecond time stamp to the 42460 millisecond timestamp is a match to the 'positive-resolution' category. In this case, the custom 'positive-resolution' category required a positive sentiment in last few seconds of speech.

```
"Categories": {
  "MatchedDetails": {
    "positive-resolution": {
      "PointsOfInterest": [
        {
          "BeginOffsetMillis": 40040,
          "EndOffsetMillis": 42460
        }
      ]
    }
  },
  "MatchedCategories": [
    "positive-resolution"
  ]
},
```

Call characteristics

Here's what call characteristics look like in your transcription output. Note that loudness scores are provided for each conversation turn, while all other characteristics are provided at the end of the transcript.

```
"LoudnessScores": [
  87.54,
  88.74,
  90.16,
  86.36,
  85.56,
  85.52,
  81.79,
  87.74,
  89.82
],
...
"ConversationCharacteristics": {
```



```

"NonTalkTime": {
  "Instances": [],
  "TotalTimeMillis": 0
},
"Interruptions": {
  "TotalCount": 2,
  "TotalTimeMillis": 10700,
  "InterruptionsByInterrupter": {
    "AGENT": [
      {
        "BeginOffsetMillis": 26040,
        "DurationMillis": 5510,
        "EndOffsetMillis": 31550
      }
    ],
    "CUSTOMER": [
      {
        "BeginOffsetMillis": 770,
        "DurationMillis": 5190,
        "EndOffsetMillis": 5960
      }
    ]
  }
},
"TotalConversationDurationMillis": 42460,
...
"TalkSpeed": {
  "DetailsByParticipant": {
    "AGENT": {
      "AverageWordsPerMinute": 150
    },
    "CUSTOMER": {
      "AverageWordsPerMinute": 167
    }
  }
},
"TalkTime": {
  "DetailsByParticipant": {
    "AGENT": {
      "TotalTimeMillis": 32750
    },
    "CUSTOMER": {
      "TotalTimeMillis": 18010
    }
  },
  "TotalTimeMillis": 50760
},
},

```

Call summarization

Here's what call summarization looks like in your transcription output:

- In the following example, **issues** are identified as starting at character 7 and ending at character 51, which refers to this section of the text: *"I would like to cancel my recipe subscription"*.

```

"Content": "Well, I would like to cancel my recipe subscription.",
"IssuesDetected": [
  {
    "CharacterOffsets": {

```

```
        "Begin": 7,  
        "End": 51  
    }  
  }  
],
```

- In the following example, **outcomes** are identified as starting at character 12 and ending at character 78, which refers to this section of the text: *"I made all changes to your account and now this discount is applied"*.

```
"Content": "Wonderful. I made all changes to your account and now this discount is  
applied, please check.",  
  
"OutcomesDetected": [  
  {  
    "CharacterOffsets": {  
      "Begin": 12,  
      "End": 78  
    }  
  }  
],
```

- In the following example, **action items** are identified as starting at character 0 and ending at character 103, which refers to this section of the text: *"I will send an email with all the details to you today, and I will call you back next week to follow up"*.

```
"Content": "I will send an email with all the details to you today, and I will call you  
back next week to follow up. Have a wonderful evening.",  
  
"ActionItemsDetected": [  
  {  
    "CharacterOffsets": {  
      "Begin": 0,  
      "End": 103  
    }  
  }  
],
```

Sensitive data redaction

Here's what sensitive data redaction looks like in your transcription output:

```
"Content": "[PII], my name is [PII], how can I help?",
```

Sentiment analysis

Here's what sentiment analysis looks like in your transcription output.

- Qualitative turn-by-turn sentiment values:

```
"Content": "That's very sad to hear. Can I offer you a 50% discount to have you stay with  
us?",  
  
...  
  
"BeginOffsetMillis": 12180,  
"EndOffsetMillis": 16960,  
"Sentiment": "NEGATIVE",
```

```
"ParticipantRole": "AGENT"
...
"Content": "That is a very generous offer. And I accept.",
...
"BeginOffsetMillis": 17140,
"EndOffsetMillis": 19860,
"Sentiment": "POSITIVE",
"ParticipantRole": "CUSTOMER"
```

- Quantitative sentiment values for the entire call:

```
"Sentiment": {
  "OverallSentiment": {
    "AGENT": 2.5,
    "CUSTOMER": 2.1
  },
}
```

- Quantitative sentiment values per participant and per call quarter:

```
"SentimentByPeriod": {
  "QUARTER": {
    "AGENT": [
      {
        "Score": 0.0,
        "BeginOffsetMillis": 0,
        "EndOffsetMillis": 9862
      },
      {
        "Score": -5.0,
        "BeginOffsetMillis": 9862,
        "EndOffsetMillis": 19725
      },
      {
        "Score": 5.0,
        "BeginOffsetMillis": 19725,
        "EndOffsetMillis": 29587
      },
      {
        "Score": 5.0,
        "BeginOffsetMillis": 29587,
        "EndOffsetMillis": 39450
      }
    ],
    "CUSTOMER": [
      {
        "Score": -2.5,
        "BeginOffsetMillis": 0,
        "EndOffsetMillis": 10615
      },
      {
        "Score": 5.0,
        "BeginOffsetMillis": 10615,
        "EndOffsetMillis": 21230
      },
      {
        "Score": 2.5,
        "BeginOffsetMillis": 21230,
        "EndOffsetMillis": 31845
      },
      {

```

```
        "Score": 5.0,  
        "BeginOffsetMillis": 31845,  
        "EndOffsetMillis": 42460  
      }  
    ]  
  }  
}
```

Compiled output

For brevity, some content is replaced with ellipses in the below transcription output.

```
{  
  "JobStatus": "COMPLETED",  
  "LanguageCode": "en-US",  
  "Transcript": [  
    {  
      "LoudnessScores": [  
        78.63,  
        78.37,  
        77.98,  
        74.18  
      ],  
      "Content": "[PII], my name is [PII], how can I help?",  
      ...  
      "Content": "Well, I would like to cancel my recipe subscription.",  
      "IssuesDetected": [  
        {  
          "CharacterOffsets": {  
            "Begin": 7,  
            "End": 51  
          }  
        }  
      ],  
      ...  
      "Content": "That's very sad to hear. Can I offer you a 50% discount to have you  
stay with us?",  
      "Items": [  
        ...  
      ],  
      "Id": "649afe93-1e59-4ae9-a3ba-a0a613868f5d",  
      "BeginOffsetMillis": 12180,  
      "EndOffsetMillis": 16960,  
      "Sentiment": "NEGATIVE",  
      "ParticipantRole": "AGENT"  
    },  
    {  
      "LoudnessScores": [  
        80.22,  
        79.48,  
        82.81  
      ],  
      "Content": "That is a very generous offer. And I accept.",  
      "Items": [  
        ...  
      ],  
      "Id": "f9266cba-34df-4ca8-9cea-4f62a52a7981",  
      "BeginOffsetMillis": 17140,  
    }  
  ]  
}
```

```
        "EndOffsetMillis": 19860,
        "Sentiment": "POSITIVE",
        "ParticipantRole": "CUSTOMER"
    },
    {
        ...

        "Content": "Wonderful. I made all changes to your account and now this discount
is applied, please check.",
        "OutcomesDetected": [
            {
                "CharacterOffsets": {
                    "Begin": 12,
                    "End": 78
                }
            }
        ],
        ...

        "Content": "I will send an email with all the details to you today, and I will
call you back next week to follow up. Have a wonderful evening.",
        "Items": [
            ...
        ],
        "Id": "78cd0923-cafd-44a5-a66e-09515796572f",
        "BeginOffsetMillis": 31800,
        "EndOffsetMillis": 39450,
        "Sentiment": "POSITIVE",
        "ParticipantRole": "AGENT"
    },
    {
        "LoudnessScores": [
            78.54,
            68.76,
            67.76
        ],
        "Content": "Thank you very much, sir. Goodbye.",
        "Items": [
            ...
        ],
        "Id": "5c5e6be0-8349-4767-8447-986f995af7c3",
        "BeginOffsetMillis": 40040,
        "EndOffsetMillis": 42460,
        "Sentiment": "POSITIVE",
        "ParticipantRole": "CUSTOMER"
    }
],
...

"Categories": {
    "MatchedDetails": {
        "positive-resolution": {
            "PointsOfInterest": [
                {
                    "BeginOffsetMillis": 40040,
                    "EndOffsetMillis": 42460
                }
            ]
        }
    }
},
"MatchedCategories": [
    "positive-resolution"
]
```

```
},  
...  
"ConversationCharacteristics": {  
  "NonTalkTime": {  
    "Instances": [],  
    "TotalTimeMillis": 0  
  },  
  "Interruptions": {  
    "TotalCount": 2,  
    "TotalTimeMillis": 10700,  
    "InterruptionsByInterrupter": {  
      "AGENT": [  
        {  
          "BeginOffsetMillis": 26040,  
          "DurationMillis": 5510,  
          "EndOffsetMillis": 31550  
        },  
        {  
          "BeginOffsetMillis": 770,  
          "DurationMillis": 5190,  
          "EndOffsetMillis": 5960  
        }  
      ]  
    }  
  },  
  "TotalConversationDurationMillis": 42460,  
  "Sentiment": {  
    "OverallSentiment": {  
      "AGENT": 2.5,  
      "CUSTOMER": 2.1  
    },  
    "SentimentByPeriod": {  
      "QUARTER": {  
        "AGENT": [  
          {  
            "Score": 0.0,  
            "BeginOffsetMillis": 0,  
            "EndOffsetMillis": 9862  
          },  
          {  
            "Score": -5.0,  
            "BeginOffsetMillis": 9862,  
            "EndOffsetMillis": 19725  
          },  
          {  
            "Score": 5.0,  
            "BeginOffsetMillis": 19725,  
            "EndOffsetMillis": 29587  
          },  
          {  
            "Score": 5.0,  
            "BeginOffsetMillis": 29587,  
            "EndOffsetMillis": 39450  
          }  
        ],  
        "CUSTOMER": [  
          {  
            "Score": -2.5,  
            "BeginOffsetMillis": 0,  
            "EndOffsetMillis": 10615  
          },  
          {
```

```
        "Score": 5.0,
        "BeginOffsetMillis": 10615,
        "EndOffsetMillis": 21230
      },
      {
        "Score": 2.5,
        "BeginOffsetMillis": 21230,
        "EndOffsetMillis": 31845
      },
      {
        "Score": 5.0,
        "BeginOffsetMillis": 31845,
        "EndOffsetMillis": 42460
      }
    ]
  }
},
"TalkSpeed": {
  "DetailsByParticipant": {
    "AGENT": {
      "AverageWordsPerMinute": 150
    },
    "CUSTOMER": {
      "AverageWordsPerMinute": 167
    }
  }
},
"TalkTime": {
  "DetailsByParticipant": {
    "AGENT": {
      "TotalTimeMillis": 32750
    },
    "CUSTOMER": {
      "TotalTimeMillis": 18010
    }
  },
  "TotalTimeMillis": 50760
}
},
```

Code examples for Amazon Transcribe using AWS SDKs

The following code examples show how to use Amazon Transcribe with an AWS software development kit (SDK).

The examples are divided into the following categories:

Actions

Code excerpts that show you how to call individual service functions.

Scenarios

Code examples that show you how to accomplish a specific task by calling multiple functions within the same service.

Cross-service examples

Sample applications that work across multiple AWS services.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Code examples

- [Actions for Amazon Transcribe using AWS SDKs \(p. 203\)](#)
 - [Create a custom Amazon Transcribe vocabulary using an AWS SDK \(p. 203\)](#)
 - [Delete a custom Amazon Transcribe vocabulary using an AWS SDK \(p. 204\)](#)
 - [Delete a Amazon Transcribe Medical transcription job using an AWS SDK \(p. 205\)](#)
 - [Delete an Amazon Transcribe transcription job using an AWS SDK \(p. 206\)](#)
 - [Get a custom Amazon Transcribe vocabulary using an AWS SDK \(p. 207\)](#)
 - [Get an Amazon Transcribe transcription job using an AWS SDK \(p. 208\)](#)
 - [List custom Amazon Transcribe vocabularies using an AWS SDK \(p. 208\)](#)
 - [List Amazon Transcribe Medical transcription jobs using an AWS SDK \(p. 209\)](#)
 - [List Amazon Transcribe transcription jobs using an AWS SDK \(p. 210\)](#)
 - [Start an Amazon Transcribe Medical transcription job using an AWS SDK \(p. 212\)](#)
 - [Start an Amazon Transcribe transcription job using an AWS SDK \(p. 213\)](#)
 - [Update a custom Amazon Transcribe vocabulary using an AWS SDK \(p. 215\)](#)
- [Scenarios for Amazon Transcribe using AWS SDKs \(p. 215\)](#)
 - [Create and refine an Amazon Transcribe custom vocabulary using an AWS SDK \(p. 216\)](#)
 - [Transcribe audio and get job data with Amazon Transcribe using an AWS SDK \(p. 221\)](#)
- [Cross-service examples for Amazon Transcribe using AWS SDKs \(p. 222\)](#)
 - [Build an Amazon Transcribe app \(p. 223\)](#)
 - [Build an Amazon Transcribe streaming app \(p. 223\)](#)

- [Convert text to speech and back to text using an AWS SDK \(p. 224\)](#)

Actions for Amazon Transcribe using AWS SDKs

The following code examples demonstrate how to perform individual Amazon Transcribe actions with AWS SDKs. These excerpts call the Amazon Transcribe API and are not intended to be run in isolation. Each example includes a link to GitHub, where you can find instructions on how to set up and run the code in context.

The following examples include only the most commonly used actions. For a complete list, see the [Amazon Transcribe API Reference](#).

Examples

- [Create a custom Amazon Transcribe vocabulary using an AWS SDK \(p. 203\)](#)
- [Delete a custom Amazon Transcribe vocabulary using an AWS SDK \(p. 204\)](#)
- [Delete a Amazon Transcribe Medical transcription job using an AWS SDK \(p. 205\)](#)
- [Delete an Amazon Transcribe transcription job using an AWS SDK \(p. 206\)](#)
- [Get a custom Amazon Transcribe vocabulary using an AWS SDK \(p. 207\)](#)
- [Get an Amazon Transcribe transcription job using an AWS SDK \(p. 208\)](#)
- [List custom Amazon Transcribe vocabularies using an AWS SDK \(p. 208\)](#)
- [List Amazon Transcribe Medical transcription jobs using an AWS SDK \(p. 209\)](#)
- [List Amazon Transcribe transcription jobs using an AWS SDK \(p. 210\)](#)
- [Start an Amazon Transcribe Medical transcription job using an AWS SDK \(p. 212\)](#)
- [Start an Amazon Transcribe transcription job using an AWS SDK \(p. 213\)](#)
- [Update a custom Amazon Transcribe vocabulary using an AWS SDK \(p. 215\)](#)

Create a custom Amazon Transcribe vocabulary using an AWS SDK

The following code example shows how to create a custom Amazon Transcribe vocabulary.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def create_vocabulary(
    vocabulary_name, language_code, transcribe_client,
    phrases=None, table_uri=None):
    """
    Creates a custom vocabulary that can be used to improve the accuracy of
    transcription jobs. This function returns as soon as the vocabulary processing
    is started. Call get_vocabulary to get the current status of the vocabulary.
    The vocabulary is ready to use when its status is 'READY'.

    :param vocabulary_name: The name of the custom vocabulary.
```

```
:param language_code: The language code of the vocabulary.  
    For example, en-US or nl-NL.  
:param transcribe_client: The Boto3 Transcribe client.  
:param phrases: A list of comma-separated phrases to include in the vocabulary.  
:param table_uri: A table of phrases and pronunciation hints to include in the  
    vocabulary.  
:return: Information about the newly created vocabulary.  
"""  
try:  
    vocab_args = {'VocabularyName': vocabulary_name, 'LanguageCode':  
language_code}  
    if phrases is not None:  
        vocab_args['Phrases'] = phrases  
    elif table_uri is not None:  
        vocab_args['VocabularyFileUri'] = table_uri  
    response = transcribe_client.create_vocabulary(**vocab_args)  
    logger.info("Created custom vocabulary %s.", response['VocabularyName'])  
except ClientError:  
    logger.exception("Couldn't create custom vocabulary %s.", vocabulary_name)  
    raise  
else:  
    return response
```

- For API details, see [CreateVocabulary](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Delete a custom Amazon Transcribe vocabulary using an AWS SDK

The following code example shows how to delete a custom Amazon Transcribe vocabulary.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def delete_vocabulary(vocabulary_name, transcribe_client):  
    """  
    Deletes a custom vocabulary.  
  
    :param vocabulary_name: The name of the vocabulary to delete.  
    :param transcribe_client: The Boto3 Transcribe client.  
    """  
    try:  
        transcribe_client.delete_vocabulary(VocabularyName=vocabulary_name)  
        logger.info("Deleted vocabulary %s.", vocabulary_name)  
    except ClientError:  
        logger.exception("Couldn't delete vocabulary %s.", vocabulary_name)  
        raise
```

- For API details, see [DeleteVocabulary](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Delete a Amazon Transcribe Medical transcription job using an AWS SDK

The following code example shows how to delete an Amazon Transcribe Medical transcription job.

JavaScript

SDK for JavaScript V3

Tip

To learn how to set up and run this example, see [GitHub](#).

Create the client.

```
const { TranscribeClient } = require("@aws-sdk/client-transcribe");
// Set the AWS Region.
const REGION = "REGION"; //e.g. "us-east-1"
// Create an Amazon Transcribe service client object.
const transcribeClient = new TranscribeClient({ region: REGION });
export { transcribeClient };
```

Delete a medical transcription job.

```
// Import the required AWS SDK clients and commands for Node.js
import { DeleteMedicalTranscriptionJobCommand } from "@aws-sdk/client-transcribe";
import { transcribeClient } from "../libs/transcribeClient.js";

// Set the parameters
export const params = {
  MedicalTranscriptionJobName: "MEDICAL_JOB_NAME", // For example,
  'medical_transcription_demo'
};

export const run = async () => {
  try {
    const data = await transcribeClient.send(
      new DeleteMedicalTranscriptionJobCommand(params)
    );
    console.log("Success - deleted");
    return data; // For unit tests.
  } catch (err) {
    console.log("Error", err);
  }
};
run();
```

- For more information, see [AWS SDK for JavaScript Developer Guide](#).
- For API details, see [DeleteMedicalTranscriptionJob](#) in *AWS SDK for JavaScript API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Delete an Amazon Transcribe transcription job using an AWS SDK

The following code examples show how to delete an Amazon Transcribe transcription job.

JavaScript

SDK for JavaScript V3

Tip

To learn how to set up and run this example, see [GitHub](#).

Create the client.

```
const { TranscribeClient } = require("@aws-sdk/client-transcribe");
// Set the AWS Region.
const REGION = "REGION"; //e.g. "us-east-1"
// Create an Amazon Transcribe service client object.
const transcribeClient = new TranscribeClient({ region: REGION });
export { transcribeClient };
```

Delete a transcription job.

```
// Import the required AWS SDK clients and commands for Node.js
import { DeleteTranscriptionJobCommand } from "@aws-sdk/client-transcribe";
import { transcribeClient } from "../libs/transcribeClient.js";

// Set the parameters
export const params = {
  TranscriptionJobName: "JOB_NAME", // Required. For example, 'transcription_demo'
};

export const run = async () => {
  try {
    const data = await transcribeClient.send(
      new DeleteTranscriptionJobCommand(params)
    );
    console.log("Success - deleted");
    return data; // For unit tests.
  } catch (err) {
    console.log("Error", err);
  }
};
run();
```

- For more information, see [AWS SDK for JavaScript Developer Guide](#).
- For API details, see [DeleteTranscriptionJob](#) in *AWS SDK for JavaScript API Reference*.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def delete_job(job_name, transcribe_client):
```

```
"""
Deletes a transcription job. This also deletes the transcript associated with
the job.

:param job_name: The name of the job to delete.
:param transcribe_client: The Boto3 Transcribe client.
"""
try:
    transcribe_client.delete_transcription_job(
        TranscriptionJobName=job_name)
    logger.info("Deleted job %s.", job_name)
except ClientError:
    logger.exception("Couldn't delete job %s.", job_name)
    raise
```

- For API details, see [DeleteTranscriptionJob](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Get a custom Amazon Transcribe vocabulary using an AWS SDK

The following code example shows how to get a custom Amazon Transcribe vocabulary.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def get_vocabulary(vocabulary_name, transcribe_client):
    """
    Gets information about a customer vocabulary.

    :param vocabulary_name: The name of the vocabulary to retrieve.
    :param transcribe_client: The Boto3 Transcribe client.
    :return: Information about the vocabulary.
    """
    try:
        response = transcribe_client.get_vocabulary(VocabularyName=vocabulary_name)
        logger.info("Got vocabulary %s.", response['VocabularyName'])
    except ClientError:
        logger.exception("Couldn't get vocabulary %s.", vocabulary_name)
        raise
    else:
        return response
```

- For API details, see [GetVocabulary](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Get an Amazon Transcribe transcription job using an AWS SDK

The following code example shows how to get an Amazon Transcribe transcription job.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def get_job(job_name, transcribe_client):
    """
    Gets details about a transcription job.

    :param job_name: The name of the job to retrieve.
    :param transcribe_client: The Boto3 Transcribe client.
    :return: The retrieved transcription job.
    """
    try:
        response = transcribe_client.get_transcription_job(
            TranscriptionJobName=job_name)
        job = response['TranscriptionJob']
        logger.info("Got job %s.", job['TranscriptionJobName'])
    except ClientError:
        logger.exception("Couldn't get job %s.", job_name)
        raise
    else:
        return job
```

- For API details, see [GetTranscriptionJob](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

List custom Amazon Transcribe vocabularies using an AWS SDK

The following code example shows how to list custom Amazon Transcribe vocabularies.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def list_vocabularies(vocabulary_filter, transcribe_client):
    """
    Lists the custom vocabularies created for this AWS account.

    :param vocabulary_filter: The returned vocabularies must contain this string in
        their names.
```

```
:param transcribe_client: The Boto3 Transcribe client.
:return: The list of retrieved vocabularies.
"""
try:
    response = transcribe_client.list_vocabularies(
        NameContains=vocabulary_filter)
    vocabs = response['Vocabularies']
    next_token = response.get('NextToken')
    while next_token is not None:
        response = transcribe_client.list_vocabularies(
            NameContains=vocabulary_filter, NextToken=next_token)
        vocabs += response['Vocabularies']
        next_token = response.get('NextToken')
    logger.info(
        "Got %s vocabularies with filter %s.", len(vocabs), vocabulary_filter)
except ClientError:
    logger.exception(
        "Couldn't list vocabularies with filter %s.", vocabulary_filter)
    raise
else:
    return vocabs
```

- For API details, see [ListVocabularies](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

List Amazon Transcribe Medical transcription jobs using an AWS SDK

The following code example shows how to list Amazon Transcribe Medical transcription jobs.

JavaScript

SDK for JavaScript V3

Tip

To learn how to set up and run this example, see [GitHub](#).

Create the client.

```
const { TranscribeClient } = require("@aws-sdk/client-transcribe");
// Set the AWS Region.
const REGION = "REGION"; //e.g. "us-east-1"
// Create an Amazon Transcribe service client object.
const transcribeClient = new TranscribeClient({ region: REGION });
export { transcribeClient };
```

List medical transcription jobs.

```
// Import the required AWS SDK clients and commands for Node.js
import { StartMedicalTranscriptionJobCommand } from "@aws-sdk/client-transcribe";
import { transcribeClient } from "../libs/transcribeClient.js";

// Set the parameters
export const params = {
    MedicalTranscriptionJobName: "MEDICAL_JOB_NAME", // Required
```

```
OutputBucketName: "OUTPUT_BUCKET_NAME", // Required
Specialty: "PRIMARYCARE", // Required. Possible values are 'PRIMARYCARE'
Type: "JOB_TYPE", // Required. Possible values are 'CONVERSATION' and 'DICTATION'
LanguageCode: "LANGUAGE_CODE", // For example, 'en-US'
MediaFormat: "SOURCE_FILE_FORMAT", // For example, 'wav'
Media: {
  MediaFileUri: "SOURCE_FILE_LOCATION",
  // The S3 object location of the input media file. The URI must be in the same
  // region
  // as the API endpoint that you are calling. For example,
  // "https://transcribe-demo.s3-REGION.amazonaws.com/hello_world.wav"
},
};

export const run = async () => {
  try {
    const data = await transcribeClient.send(
      new StartMedicalTranscriptionJobCommand(params)
    );
    console.log("Success - put", data);
    return data; // For unit tests.
  } catch (err) {
    console.log("Error", err);
  }
};
run();
```

- For more information, see [AWS SDK for JavaScript Developer Guide](#).
- For API details, see [ListMedicalTranscriptionJobs](#) in *AWS SDK for JavaScript API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

List Amazon Transcribe transcription jobs using an AWS SDK

The following code examples show how to list Amazon Transcribe transcription jobs.

JavaScript

SDK for JavaScript V3

Tip

To learn how to set up and run this example, see [GitHub](#).

Create the client.

```
const { TranscribeClient } = require("@aws-sdk/client-transcribe");
// Set the AWS Region.
const REGION = "REGION"; //e.g. "us-east-1"
// Create an Amazon Transcribe service client object.
const transcribeClient = new TranscribeClient({ region: REGION });
export { transcribeClient };
```

List transcription jobs.

```
// Import the required AWS SDK clients and commands for Node.js
```



```
import { ListTranscriptionJobsCommand } from "@aws-sdk/client-transcribe";
import { transcribeClient } from "../libs/transcribeClient.js";

// Set the parameters
export const params = {
  JobNameContains: "KEYWORD", // Not required. Returns only transcription
  // job names containing this string
};

export const run = async () => {
  try {
    const data = await transcribeClient.send(
      new ListTranscriptionJobsCommand(params)
    );
    console.log("Success", data.TranscriptionJobSummaries);
    return data; // For unit tests.
  } catch (err) {
    console.log("Error", err);
  }
};
run();
```

- For more information, see [AWS SDK for JavaScript Developer Guide](#).
- For API details, see [ListTranscriptionJobs](#) in *AWS SDK for JavaScript API Reference*.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def list_jobs(job_filter, transcribe_client):
    """
    Lists summaries of the transcription jobs for the current AWS account.

    :param job_filter: The list of returned jobs must contain this string in their
        names.
    :param transcribe_client: The Boto3 Transcribe client.
    :return: The list of retrieved transcription job summaries.
    """
    try:
        response = transcribe_client.list_transcription_jobs(
            JobNameContains=job_filter)
        jobs = response['TranscriptionJobSummaries']
        next_token = response.get('NextToken')
        while next_token is not None:
            response = transcribe_client.list_transcription_jobs(
                JobNameContains=job_filter, NextToken=next_token)
            jobs += response['TranscriptionJobSummaries']
            next_token = response.get('NextToken')
        logger.info("Got %s jobs with filter %s.", len(jobs), job_filter)
    except ClientError:
        logger.exception("Couldn't get jobs with filter %s.", job_filter)
        raise
    else:
        return jobs
```

- For API details, see [ListTranscriptionJobs](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Start an Amazon Transcribe Medical transcription job using an AWS SDK

The following code example shows how to start an Amazon Transcribe Medical transcription job.

JavaScript

SDK for JavaScript V3

Tip

To learn how to set up and run this example, see [GitHub](#).

Create the client.

```
const { TranscribeClient } = require("@aws-sdk/client-transcribe");
// Set the AWS Region.
const REGION = "REGION"; //e.g. "us-east-1"
// Create an Amazon Transcribe service client object.
const transcribeClient = new TranscribeClient({ region: REGION });
export { transcribeClient };
```

Start a medical transcription job.

```
// Import the required AWS SDK clients and commands for Node.js
import { StartMedicalTranscriptionJobCommand } from "@aws-sdk/client-transcribe";
import { transcribeClient } from "../libs/transcribeClient.js";

// Set the parameters
export const params = {
  MedicalTranscriptionJobName: "MEDICAL_JOB_NAME", // Required
  OutputBucketName: "OUTPUT_BUCKET_NAME", // Required
  Specialty: "PRIMARYCARE", // Required. Possible values are 'PRIMARYCARE'
  Type: "JOB_TYPE", // Required. Possible values are 'CONVERSATION' and 'DICTATION'
  LanguageCode: "LANGUAGE_CODE", // For example, 'en-US'
  MediaFormat: "SOURCE_FILE_FORMAT", // For example, 'wav'
  Media: {
    MediaFileUri: "SOURCE_FILE_LOCATION",
    // The S3 object location of the input media file. The URI must be in the same
    // region
    // as the API endpoint that you are calling. For example,
    // "https://transcribe-demo.s3-REGION.amazonaws.com/hello_world.wav"
  },
};

export const run = async () => {
  try {
    const data = await transcribeClient.send(
      new StartMedicalTranscriptionJobCommand(params)
    );
    console.log("Success - put", data);
    return data; // For unit tests.
  } catch (err) {
    console.log("Error", err);
  }
};
run();
```

- For more information, see [AWS SDK for JavaScript Developer Guide](#).
- For API details, see [StartMedicalTranscriptionJob](#) in *AWS SDK for JavaScript API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Start an Amazon Transcribe transcription job using an AWS SDK

The following code examples show how to start an Amazon Transcribe transcription job.

JavaScript

SDK for JavaScript V3

Tip

To learn how to set up and run this example, see [GitHub](#).

Create the client.

```
const { TranscribeClient } = require("@aws-sdk/client-transcribe");
// Set the AWS Region.
const REGION = "REGION"; //e.g. "us-east-1"
// Create an Amazon Transcribe service client object.
const transcribeClient = new TranscribeClient({ region: REGION });
export { transcribeClient };
```

Start a transcription job.

```
// Import the required AWS SDK clients and commands for Node.js
import { StartTranscriptionJobCommand } from "@aws-sdk/client-transcribe";
import { transcribeClient } from "../libs/transcribeClient.js";

// Set the parameters
export const params = {
  TranscriptionJobName: "JOB_NAME",
  LanguageCode: "LANGUAGE_CODE", // For example, 'en-US'
  MediaFormat: "SOURCE_FILE_FORMAT", // For example, 'wav'
  Media: {
    MediaFileUri: "SOURCE_LOCATION",
    // For example, "https://transcribe-demo.s3-REGION.amazonaws.com/
hello_world.wav"
  },
  OutputBucketName: "OUTPUT_BUCKET_NAME"
};

export const run = async () => {
  try {
    const data = await transcribeClient.send(
      new StartTranscriptionJobCommand(params)
    );
    console.log("Success - put", data);
    return data; // For unit tests.
  } catch (err) {
    console.log("Error", err);
  }
};
```

```
    }  
};  
run();
```

- For more information, see [AWS SDK for JavaScript Developer Guide](#).
- For API details, see [StartTranscriptionJob](#) in *AWS SDK for JavaScript API Reference*.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def start_job(  
    job_name, media_uri, media_format, language_code, transcribe_client,  
    vocabulary_name=None):  
    """  
    Starts a transcription job. This function returns as soon as the job is  
    started.  
    To get the current status of the job, call get_transcription_job. The job is  
    successfully completed when the job status is 'COMPLETED'.  
  
    :param job_name: The name of the transcription job. This must be unique for  
        your AWS account.  
    :param media_uri: The URI where the audio file is stored. This is typically  
        in an Amazon S3 bucket.  
    :param media_format: The format of the audio file. For example, mp3 or wav.  
    :param language_code: The language code of the audio file.  
        For example, en-US or ja-JP  
    :param transcribe_client: The Boto3 Transcribe client.  
    :param vocabulary_name: The name of a custom vocabulary to use when  
    transcribing  
        the audio file.  
    :return: Data about the job.  
    """  
    try:  
        job_args = {  
            'TranscriptionJobName': job_name,  
            'Media': {'MediaFileUri': media_uri},  
            'MediaFormat': media_format,  
            'LanguageCode': language_code}  
        if vocabulary_name is not None:  
            job_args['Settings'] = {'VocabularyName': vocabulary_name}  
        response = transcribe_client.start_transcription_job(**job_args)  
        job = response['TranscriptionJob']  
        logger.info("Started transcription job %s.", job_name)  
    except ClientError:  
        logger.exception("Couldn't start transcription job %s.", job_name)  
        raise  
    else:  
        return job
```

- For API details, see [StartTranscriptionJob](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Update a custom Amazon Transcribe vocabulary using an AWS SDK

The following code example shows how to update a custom Amazon Transcribe vocabulary.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
def update_vocabulary(
    vocabulary_name, language_code, transcribe_client, phrases=None,
    table_uri=None):
    """
    Updates an existing custom vocabulary. The entire vocabulary is replaced with
    the contents of the update.

    :param vocabulary_name: The name of the vocabulary to update.
    :param language_code: The language code of the vocabulary.
    :param transcribe_client: The Boto3 Transcribe client.
    :param phrases: A list of comma-separated phrases to include in the vocabulary.
    :param table_uri: A table of phrases and pronunciation hints to include in the
        vocabulary.
    """
    try:
        vocab_args = {'VocabularyName': vocabulary_name, 'LanguageCode':
language_code}
        if phrases is not None:
            vocab_args['Phrases'] = phrases
        elif table_uri is not None:
            vocab_args['VocabularyFileUri'] = table_uri
        response = transcribe_client.update_vocabulary(**vocab_args)
        logger.info(
            "Updated custom vocabulary %s.", response['VocabularyName'])
    except ClientError:
        logger.exception("Couldn't update custom vocabulary %s.", vocabulary_name)
        raise
```

- For API details, see [UpdateVocabulary](#) in *AWS SDK for Python (Boto3) API Reference*.

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Scenarios for Amazon Transcribe using AWS SDKs

The following code examples show you how to implement common scenarios in Amazon Transcribe with AWS SDKs. These scenarios show you how to accomplish specific tasks by calling multiple functions within Amazon Transcribe. Each scenario includes a link to GitHub, where you can find instructions on how to set up and run the code.

Examples

- [Create and refine an Amazon Transcribe custom vocabulary using an AWS SDK \(p. 216\)](#)

- [Transcribe audio and get job data with Amazon Transcribe using an AWS SDK \(p. 221\)](#)

Create and refine an Amazon Transcribe custom vocabulary using an AWS SDK

The following code example shows how to:

- Upload an audio file to Amazon S3.
- Run an Amazon Transcribe job to transcribe the file and get the results.
- Create and refine a custom vocabulary to improve transcription accuracy.
- Run jobs with custom vocabularies and get the results.

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

Transcribe an audio file that contains a reading of Jabberwocky by Lewis Carroll. Start by creating functions that wrap Amazon Transcribe actions.

```
def start_job(
    job_name, media_uri, media_format, language_code, transcribe_client,
    vocabulary_name=None):
    """
    Starts a transcription job. This function returns as soon as the job is
    started.
    To get the current status of the job, call get_transcription_job. The job is
    successfully completed when the job status is 'COMPLETED'.

    :param job_name: The name of the transcription job. This must be unique for
        your AWS account.
    :param media_uri: The URI where the audio file is stored. This is typically
        in an Amazon S3 bucket.
    :param media_format: The format of the audio file. For example, mp3 or wav.
    :param language_code: The language code of the audio file.
        For example, en-US or ja-JP
    :param transcribe_client: The Boto3 Transcribe client.
    :param vocabulary_name: The name of a custom vocabulary to use when
    transcribing
        the audio file.
    :return: Data about the job.
    """
    try:
        job_args = {
            'TranscriptionJobName': job_name,
            'Media': {'MediaFileUri': media_uri},
            'MediaFormat': media_format,
            'LanguageCode': language_code}
        if vocabulary_name is not None:
            job_args['Settings'] = {'VocabularyName': vocabulary_name}
        response = transcribe_client.start_transcription_job(**job_args)
        job = response['TranscriptionJob']
        logger.info("Started transcription job %s.", job_name)
    except ClientError:
        logger.exception("Couldn't start transcription job %s.", job_name)
        raise
    else:
```

```
        return job

def get_job(job_name, transcribe_client):
    """
    Gets details about a transcription job.

    :param job_name: The name of the job to retrieve.
    :param transcribe_client: The Boto3 Transcribe client.
    :return: The retrieved transcription job.
    """
    try:
        response = transcribe_client.get_transcription_job(
            TranscriptionJobName=job_name)
        job = response['TranscriptionJob']
        logger.info("Got job %s.", job['TranscriptionJobName'])
    except ClientError:
        logger.exception("Couldn't get job %s.", job_name)
        raise
    else:
        return job

def delete_job(job_name, transcribe_client):
    """
    Deletes a transcription job. This also deletes the transcript associated with
    the job.

    :param job_name: The name of the job to delete.
    :param transcribe_client: The Boto3 Transcribe client.
    """
    try:
        transcribe_client.delete_transcription_job(
            TranscriptionJobName=job_name)
        logger.info("Deleted job %s.", job_name)
    except ClientError:
        logger.exception("Couldn't delete job %s.", job_name)
        raise

def create_vocabulary(
    vocabulary_name, language_code, transcribe_client,
    phrases=None, table_uri=None):
    """
    Creates a custom vocabulary that can be used to improve the accuracy of
    transcription jobs. This function returns as soon as the vocabulary processing
    is started. Call get_vocabulary to get the current status of the vocabulary.
    The vocabulary is ready to use when its status is 'READY'.

    :param vocabulary_name: The name of the custom vocabulary.
    :param language_code: The language code of the vocabulary.
        For example, en-US or nl-NL.
    :param transcribe_client: The Boto3 Transcribe client.
    :param phrases: A list of comma-separated phrases to include in the vocabulary.
    :param table_uri: A table of phrases and pronunciation hints to include in the
        vocabulary.
    :return: Information about the newly created vocabulary.
    """
    try:
        vocab_args = {'VocabularyName': vocabulary_name, 'LanguageCode':
language_code}
        if phrases is not None:
            vocab_args['Phrases'] = phrases
        elif table_uri is not None:
            vocab_args['VocabularyFileUri'] = table_uri
        response = transcribe_client.create_vocabulary(**vocab_args)
        logger.info("Created custom vocabulary %s.", response['VocabularyName'])
    except ClientError:
        logger.exception("Couldn't create custom vocabulary %s.", vocabulary_name)
```

```
        raise
    else:
        return response

def get_vocabulary(vocabulary_name, transcribe_client):
    """
    Gets information about a customer vocabulary.

    :param vocabulary_name: The name of the vocabulary to retrieve.
    :param transcribe_client: The Boto3 Transcribe client.
    :return: Information about the vocabulary.
    """
    try:
        response = transcribe_client.get_vocabulary(VocabularyName=vocabulary_name)
        logger.info("Got vocabulary %s.", response['VocabularyName'])
    except ClientError:
        logger.exception("Couldn't get vocabulary %s.", vocabulary_name)
        raise
    else:
        return response

def update_vocabulary(
    vocabulary_name, language_code, transcribe_client, phrases=None,
    table_uri=None):
    """
    Updates an existing custom vocabulary. The entire vocabulary is replaced with
    the contents of the update.

    :param vocabulary_name: The name of the vocabulary to update.
    :param language_code: The language code of the vocabulary.
    :param transcribe_client: The Boto3 Transcribe client.
    :param phrases: A list of comma-separated phrases to include in the vocabulary.
    :param table_uri: A table of phrases and pronunciation hints to include in the
        vocabulary.
    """
    try:
        vocab_args = {'VocabularyName': vocabulary_name, 'LanguageCode':
language_code}
        if phrases is not None:
            vocab_args['Phrases'] = phrases
        elif table_uri is not None:
            vocab_args['VocabularyFileUri'] = table_uri
        response = transcribe_client.update_vocabulary(**vocab_args)
        logger.info(
            "Updated custom vocabulary %s.", response['VocabularyName'])
    except ClientError:
        logger.exception("Couldn't update custom vocabulary %s.", vocabulary_name)
        raise

def list_vocabularies(vocabulary_filter, transcribe_client):
    """
    Lists the custom vocabularies created for this AWS account.

    :param vocabulary_filter: The returned vocabularies must contain this string in
        their names.
    :param transcribe_client: The Boto3 Transcribe client.
    :return: The list of retrieved vocabularies.
    """
    try:
        response = transcribe_client.list_vocabularies(
            NameContains=vocabulary_filter)
        vocabs = response['Vocabularies']
        next_token = response.get('NextToken')
        while next_token is not None:
            response = transcribe_client.list_vocabularies(
                NameContains=vocabulary_filter, NextToken=next_token)
```



```
        vocabs += response['Vocabularies']
        next_token = response.get('NextToken')
    logger.info(
        "Got %s vocabularies with filter %s.", len(vocabs), vocabulary_filter)
except ClientError:
    logger.exception(
        "Couldn't list vocabularies with filter %s.", vocabulary_filter)
    raise
else:
    return vocabs

def delete_vocabulary(vocabulary_name, transcribe_client):
    """
    Deletes a custom vocabulary.

    :param vocabulary_name: The name of the vocabulary to delete.
    :param transcribe_client: The Boto3 Transcribe client.
    """
    try:
        transcribe_client.delete_vocabulary(VocabularyName=vocabulary_name)
        logger.info("Deleted vocabulary %s.", vocabulary_name)
    except ClientError:
        logger.exception("Couldn't delete vocabulary %s.", vocabulary_name)
        raise
```

Call the wrapper functions to transcribe audio without a custom vocabulary and then with different versions of a custom vocabulary to see improved results.

```
def usage_demo():
    """Shows how to use the Amazon Transcribe service."""
    logging.basicConfig(level=logging.INFO, format='%(levelname)s: %(message)s')

    s3_resource = boto3.resource('s3')
    transcribe_client = boto3.client('transcribe')

    print('-'*88)
    print("Welcome to the Amazon Transcribe demo!")
    print('-'*88)

    bucket_name = f'jabber-bucket-{time.time_ns()}'
    print(f"Creating bucket {bucket_name}.")
    bucket = s3_resource.create_bucket(
        Bucket=bucket_name,
        CreateBucketConfiguration={
            'LocationConstraint': transcribe_client.meta.region_name})
    media_file_name = '.media/Jabberwocky.mp3'
    media_object_key = 'Jabberwocky.mp3'
    print(f"Uploading media file {media_file_name}.")
    bucket.upload_file(media_file_name, media_object_key)
    media_uri = f's3://{bucket.name}/{media_object_key}'

    job_name_simple = f'Jabber-{time.time_ns()}'
    print(f"Starting transcription job {job_name_simple}.")
    start_job(
        job_name_simple, f's3://{bucket.name}/{media_object_key}', 'mp3', 'en-US',
        transcribe_client)
    transcribe_waiter = TranscribeCompleteWaiter(transcribe_client)
    transcribe_waiter.wait(job_name_simple)
    job_simple = get_job(job_name_simple, transcribe_client)
    transcript_simple = requests.get(
        job_simple['Transcript']['TranscriptFileUri']).json()
    print(f"Transcript for job {transcript_simple['jobName']}:")
    print(transcript_simple['results'][0]['transcripts'][0]['transcript'])
```

```
print('-'*88)
print("Creating a custom vocabulary that lists the nonsense words to try to "
      "improve the transcription.")
vocabulary_name = f'Jabber-vocabulary-{time.time_ns()}'
create_vocabulary(
    vocabulary_name, 'en-US', transcribe_client,
    phrases=[
        'brillig', 'slithy', 'borogoves', 'mome', 'raths', 'Jub-Jub',
'frumious',
        'manxome', 'Tumtum', 'uffish', 'whiffling', 'tulgey', 'thou',
'frabjous',
        'callooh', 'callay', 'chortled'],
    )
vocabulary_ready_waiter = VocabularyReadyWaiter(transcribe_client)
vocabulary_ready_waiter.wait(vocabulary_name)

job_name_vocabulary_list = f'Jabber-vocabulary-list-{time.time_ns()}'
print(f"Starting transcription job {job_name_vocabulary_list}.")
start_job(
    job_name_vocabulary_list, media_uri, 'mp3', 'en-US', transcribe_client,
    vocabulary_name)
transcribe_waiter.wait(job_name_vocabulary_list)
job_vocabulary_list = get_job(job_name_vocabulary_list, transcribe_client)
transcript_vocabulary_list = requests.get(
    job_vocabulary_list['Transcript']['TranscriptFileUri']).json()
print(f"Transcript for job {transcript_vocabulary_list['jobName']}:")
print(transcript_vocabulary_list['results']['transcripts'][0]['transcript'])

print('-'*88)
print("Updating the custom vocabulary with table data that provides additional
"
      "pronunciation hints.")
table_vocab_file = 'jabber-vocabulary-table.txt'
bucket.upload_file(table_vocab_file, table_vocab_file)
update_vocabulary(
    vocabulary_name, 'en-US', transcribe_client,
    table_uri=f's3://{bucket.name}/{table_vocab_file}')
vocabulary_ready_waiter.wait(vocabulary_name)

job_name_vocab_table = f'Jabber-vocab-table-{time.time_ns()}'
print(f"Starting transcription job {job_name_vocab_table}.")
start_job(
    job_name_vocab_table, media_uri, 'mp3', 'en-US', transcribe_client,
    vocabulary_name=vocabulary_name)
transcribe_waiter.wait(job_name_vocab_table)
job_vocab_table = get_job(job_name_vocab_table, transcribe_client)
transcript_vocab_table = requests.get(
    job_vocab_table['Transcript']['TranscriptFileUri']).json()
print(f"Transcript for job {transcript_vocab_table['jobName']}:")
print(transcript_vocab_table['results']['transcripts'][0]['transcript'])

print('-'*88)
print("Getting data for jobs and vocabularies.")
jabber_jobs = list_jobs('Jabber', transcribe_client)
print(f"Found {len(jabber_jobs)} jobs:")
for job_sum in jabber_jobs:
    job = get_job(job_sum['TranscriptionJobName'], transcribe_client)
    print(f"\t{job['TranscriptionJobName']}, {job['Media']['MediaFileUri']}, "
          f"{job['Settings'].get('VocabularyName')}")

jabber_vocabs = list_vocabularies('Jabber', transcribe_client)
print(f"Found {len(jabber_vocabs)} vocabularies:")
for vocab_sum in jabber_vocabs:
    vocab = get_vocabulary(vocab_sum['VocabularyName'], transcribe_client)
    vocab_content = requests.get(vocab['DownloadUri']).text
```

```
print(f"\t{vocab['VocabularyName']} contents:")
print(vocab_content)

print('-'*88)
print("Deleting demo jobs.")
for job_name in [job_name_simple, job_name_vocabulary_list,
job_name_vocab_table]:
    delete_job(job_name, transcribe_client)
print("Deleting demo vocabulary.")
delete_vocabulary(vocabulary_name, transcribe_client)
print("Deleting demo bucket.")
bucket.objects.delete()
bucket.delete()
print("Thanks for watching!")
```

- For API details, see the following topics in *AWS SDK for Python (Boto3) API Reference*.
 - [CreateVocabulary](#)
 - [DeleteTranscriptionJob](#)
 - [DeleteVocabulary](#)
 - [GetTranscriptionJob](#)
 - [GetVocabulary](#)
 - [ListVocabularies](#)
 - [StartTranscriptionJob](#)
 - [UpdateVocabulary](#)

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Transcribe audio and get job data with Amazon Transcribe using an AWS SDK

The following code example shows how to:

- Start a transcription job with Amazon Transcribe.
- Wait for the job to complete.
- Get the URI where the transcript is stored.

For more information, see [Getting started with Amazon Transcribe](#).

Python

SDK for Python (Boto3)

Tip

To learn how to set up and run this example, see [GitHub](#).

```
import time
import boto3

def transcribe_file(job_name, file_uri, transcribe_client):
```

```
transcribe_client.start_transcription_job(
    TranscriptionJobName=job_name,
    Media={'MediaFileUri': file_uri},
    MediaFormat='wav',
    LanguageCode='en-US'
)

max_tries = 60
while max_tries > 0:
    max_tries -= 1
    job =
transcribe_client.get_transcription_job(TranscriptionJobName=job_name)
    job_status = job['TranscriptionJob']['TranscriptionJobStatus']
    if job_status in ['COMPLETED', 'FAILED']:
        print(f"Job {job_name} is {job_status}.")
        if job_status == 'COMPLETED':
            print(
                f"Download the transcript from\n"
                f"\t{job['TranscriptionJob']['Transcript']
['TranscriptFileUri']}.")
            break
        else:
            print(f"Waiting for {job_name}. Current status is {job_status}.")
            time.sleep(10)

def main():
    transcribe_client = boto3.client('transcribe')
    file_uri = 's3://test-transcribe/answer2.wav'
    transcribe_file('Example-job', file_uri, transcribe_client)

if __name__ == '__main__':
    main()
```

- For API details, see the following topics in *AWS SDK for Python (Boto3) API Reference*.
 - [GetTranscriptionJob](#)
 - [StartTranscriptionJob](#)

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Cross-service examples for Amazon Transcribe using AWS SDKs

The following sample applications use AWS SDKs to combine Amazon Transcribe with other AWS services. Each example includes a link to GitHub, where you can find instructions on how to set up and run the application.

Examples

- [Build an Amazon Transcribe app \(p. 223\)](#)
- [Build an Amazon Transcribe streaming app \(p. 223\)](#)
- [Convert text to speech and back to text using an AWS SDK \(p. 224\)](#)

Build an Amazon Transcribe app

The following code example shows how to use Amazon Transcribe to transcribe and display voice recordings in the browser.

JavaScript

SDK for JavaScript V3

Create an app that uses Amazon Transcribe to transcribe and display voice recordings in the browser. The app uses two Amazon Simple Storage Service (Amazon S3) buckets, one to host the application code, and another to store transcriptions. The app uses an Amazon Cognito user pool to authenticate your users. Authenticated users have AWS Identity and Access Management (IAM) permissions to access the required AWS services.

For complete source code and instructions on how to set up and run, see the full example on [GitHub](#).

This example is also available in the [AWS SDK for JavaScript v3 developer guide](#).

Services used in this example

- Amazon Cognito Identity
- Amazon S3
- Amazon Transcribe

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Build an Amazon Transcribe streaming app

The following code example shows how to build an app that records, transcribes, and translates live audio in real-time, and emails the results.

JavaScript

SDK for JavaScript V3

Shows how to use Amazon Transcribe to build an app that records, transcribes, and translates live audio in real-time, and emails the results using Amazon Simple Email Service (Amazon SES).

For complete source code and instructions on how to set up and run, see the full example on [GitHub](#).

Services used in this example

- Amazon Comprehend
- Amazon SES
- Amazon Transcribe
- Amazon Translate

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Convert text to speech and back to text using an AWS SDK

The following code example shows how to:

- Use Amazon Polly to synthesize a plain text (UTF-8) input file to an audio file.
- Upload the audio file to an Amazon Simple Storage Service (Amazon S3) bucket.
- Use Amazon Transcribe to convert the audio file to text.
- Display the text.

Rust

SDK for Rust

Note

This documentation is for an SDK in preview release. The SDK is subject to change and should not be used in production.

Use Amazon Polly to synthesize a plain text (UTF-8) input file to an audio file, upload the audio file to an Amazon Simple Storage Service bucket, use Amazon Transcribe to convert that audio file to text, and display the text.

For complete source code and instructions on how to set up and run, see the full example on [GitHub](#).

Services used in this example

- Amazon Polly
- Amazon S3
- Amazon Transcribe

For a complete list of AWS SDK developer guides and code examples, see [Using this service with an AWS SDK \(p. 74\)](#). This topic also includes information about getting started and details about previous SDK versions.

Security in Amazon Transcribe

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture 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 you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to Amazon Transcribe, 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 company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon Transcribe. The following topics show you how to configure Amazon Transcribe to meet your security and compliance objectives. You also learn how to use other AWS services to monitor and secure your Amazon Transcribe resources.

Topics

- [Identity and Access Management for Amazon Transcribe \(p. 225\)](#)
- [Data protection in Amazon Transcribe \(p. 242\)](#)
- [Monitoring Amazon Transcribe \(p. 245\)](#)
- [Compliance Validation for Amazon Transcribe \(p. 252\)](#)
- [Resilience in Amazon Transcribe \(p. 252\)](#)
- [Infrastructure Security in Amazon Transcribe \(p. 253\)](#)
- [Vulnerability analysis and management in Amazon Transcribe \(p. 253\)](#)
- [Security Best Practices for Amazon Transcribe \(p. 255\)](#)

Identity and Access Management for Amazon Transcribe

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 Transcribe resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience \(p. 226\)](#)
- [Authenticating with identities \(p. 226\)](#)
- [Managing access using policies \(p. 228\)](#)
- [How Amazon Transcribe works with IAM \(p. 229\)](#)
- [Cross-service confused deputy prevention \(p. 234\)](#)
- [Amazon Transcribe identity-based policy examples \(p. 235\)](#)

- [Troubleshooting Amazon Transcribe identity and access \(p. 240\)](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Amazon Transcribe.

Service user – If you use the Amazon Transcribe service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon Transcribe 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 Transcribe, see [Troubleshooting Amazon Transcribe identity and access \(p. 240\)](#).

Service administrator – If you're in charge of Amazon Transcribe resources at your company, you probably have full access to Amazon Transcribe. It's your job to determine which Amazon Transcribe features and resources your employees 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 Transcribe, see [How Amazon Transcribe works with IAM \(p. 229\)](#).

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 Transcribe. To view example Amazon Transcribe identity-based policies that you can use in IAM, see [Amazon Transcribe identity-based policy examples \(p. 235\)](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. For more information about signing in using the AWS Management Console, see [Signing in to the AWS Management Console as an IAM user or root user](#) in the *IAM User Guide*.

You must be *authenticated* (signed in to AWS) as the AWS account root user, an IAM user, or by assuming an IAM role. You can also use your company's single sign-on authentication or even sign in using Google or Facebook. In these cases, your administrator previously set up identity federation using IAM roles. When you access AWS using credentials from another company, you are assuming a role indirectly.

To sign in directly to the [AWS Management Console](#), use your password with your root user email address or your IAM user name. You can access AWS programmatically using your root user or IAM users access keys. AWS provides SDK and command line tools to cryptographically sign your request using your credentials. If you don't use AWS tools, you must sign the request yourself. Do this using *Signature Version 4*, a protocol for authenticating inbound API requests. For more information about authenticating requests, see [Signature Version 4 signing process](#) in the *AWS General Reference*.

Regardless of the authentication method that you use, you might also 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 [Using multi-factor authentication \(MFA\) in AWS](#) in the *IAM User Guide*.

AWS account root user

When you first create an AWS account, you begin with a single 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 do not use the root user for your everyday tasks, even the administrative ones. Instead, adhere to the [best practice of using the root user only to create your first IAM user](#). Then securely lock away the root user credentials and use them to perform only a few account and service management tasks.

IAM users and groups

An *IAM user* is an identity within your AWS account that has specific permissions for a single person or application. An IAM user can have long-term credentials such as a user name and password or a set of access keys. To learn how to generate access keys, see [Managing access keys for IAM users](#) in the *IAM User Guide*. When you generate access keys for an IAM user, make sure you view and securely save the key pair. You cannot recover the secret access key in the future. Instead, you must generate a new access key pair.

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 [When to create an IAM user \(instead of a role\)](#) 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. You can temporarily assume an IAM role in the AWS Management Console by [switching roles](#). 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 [Using IAM roles](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Temporary IAM user permissions** – An IAM user can assume an IAM role to temporarily take on different permissions for a specific task.
- **Federated user access** – Instead of creating an IAM user, you can use existing identities from AWS Directory Service, your enterprise user directory, or a web identity provider. These are known as *federated users*. AWS assigns a role to a federated user when access is requested through an [identity provider](#). For more information about federated users, see [Federated users and roles](#) in the *IAM User Guide*.
- **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 [How IAM roles differ from resource-based policies](#) 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.
- **Principal permissions** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. Policies grant permissions to a principal. When you use some services, you might perform an action that then triggers another action in a different service. In this case, you must have permissions to perform both actions. To see whether an action requires additional dependent actions in a policy, see [Actions, resources, and condition keys for Amazon Transcribe](#) in the *Service Authorization Reference*.
- **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 [Creating 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 IAM 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 [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see [When to create an IAM role \(instead of a user\)](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to IAM identities or AWS resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. You can sign in as the root user or an IAM user, or you can assume an IAM role. When you then make a request, AWS evaluates the related identity-based or resource-based policies. 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**.

Every IAM entity (user or role) starts with no permissions. In other words, by default, users can do nothing, not even change their own password. To give a user permission to do something, an administrator must attach a permissions policy to a user. Or the administrator can add the user to a group that has the intended permissions. When an administrator gives permissions to a group, all users in that group are granted those permissions.

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 [Creating IAM 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 [Choosing 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 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 [How SCPs work](#) 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 Transcribe works with IAM

Before you use IAM to manage access to Amazon Transcribe, learn what IAM features are available to use with Amazon Transcribe.

IAM features you can use with Amazon Transcribe

IAM feature	Amazon Transcribe support
Identity-based policies (p. 230)	Yes
Resource-based policies (p. 230)	No

IAM feature	Amazon Transcribe support
Policy actions (p. 231)	Yes
Policy resources (p. 231)	Yes
Policy condition keys (service-specific) (p. 232)	Yes
ACLs (p. 232)	No
ABAC (tags in policies) (p. 233)	Partial
Temporary credentials (p. 233)	Yes
Principal permissions (p. 233)	Yes
Service roles (p. 234)	Yes
Service-linked roles (p. 234)	No

To get a high-level view of how Amazon Transcribe 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 Transcribe

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 [Creating IAM 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 Transcribe

To view examples of Amazon Transcribe identity-based policies, see [Amazon Transcribe identity-based policy examples \(p. 235\)](#).

Resource-based policies within Amazon Transcribe

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 [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.

Policy actions for Amazon Transcribe

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 Transcribe actions, see [Actions defined by Amazon Transcribe](#) in the *Service Authorization Reference*.

Policy actions in Amazon Transcribe use the `transcribe` prefix before the action. To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
  "transcribe:action1",
  "transcribe:action2"
]
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `List`, include the following action:

```
"Action": "transcribe:List*"
```

To view examples of Amazon Transcribe identity-based policies, see [Amazon Transcribe identity-based policy examples](#) (p. 235).

Policy resources for Amazon Transcribe

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": "*" 
```

To see a list of Amazon Transcribe resource types and their ARNs, see [Resources defined by Amazon Transcribe](#) in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see [Actions defined by Amazon Transcribe](#).

To view examples of Amazon Transcribe identity-based policies, see [Amazon Transcribe identity-based policy examples \(p. 235\)](#).

Policy condition keys for Amazon Transcribe

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 Transcribe condition keys, see [Condition keys for Amazon Transcribe](#) in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions defined by Amazon Transcribe](#).

To view examples of Amazon Transcribe identity-based policies, see [Amazon Transcribe identity-based policy examples \(p. 235\)](#).

ACLs in Amazon Transcribe

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 Transcribe

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.

For more information about ABAC, see [What is ABAC?](#) 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*.

For more information about tagging Amazon Transcribe resources, see [Tagging resources \(p. 96\)](#). For more detailed information on tag-based access control, see [Controlling access to AWS resources using tags](#).

Using temporary credentials with Amazon Transcribe

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 [Switching to a 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](#).

Cross-service principal permissions for Amazon Transcribe

Supports principal permissions	Yes
--------------------------------	-----

When you use an IAM user or role to perform actions in AWS, you are considered a principal. Policies grant permissions to a principal. When you use some services, you might perform an action that then triggers another action in a different service. In this case, you must have permissions to perform both actions. To see whether an action requires additional dependent actions in a policy, see [Actions, resources, and condition keys for Amazon Transcribe](#) in the *Service Authorization Reference*.

Service roles for Amazon Transcribe

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 [Creating 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 Transcribe functionality. Edit service roles only when Amazon Transcribe provides guidance to do so.

Service-linked roles for Amazon Transcribe

Supports service-linked roles	No
-------------------------------	----

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 IAM account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

Amazon Transcribe doesn't support service-linked roles.

For details about creating or managing service-linked roles for other services, see [AWS services that work with IAM](#). Find a service in the table that includes a **Yes** in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Cross-service confused deputy prevention

A confused deputy is an entity (a service or an account) that is coerced by a different entity to perform an action. This type of impersonation can happen cross-account and cross-service.

To prevent confused deputies, AWS provides tools that help you protect your data for all services using service principals that have been given access to resources in your AWS account. This section focuses on cross-service confused deputy prevention specific to Amazon Transcribe; however, you can learn more about this topic in the [confused deputy problem](#) section of the *IAM User Guide*.

To limit the permissions IAM gives to Amazon Transcribe to access your resources, we recommend using the global condition context keys `aws:SourceArn` and `aws:SourceAccount` in your resource policies. Note that if you use both of these global condition context keys, and the `aws:SourceArn` value contains the AWS account ID, the `aws:SourceAccount` value and the AWS account in the `aws:SourceArn` value must use the same AWS account ID when used in the same policy statement.

If you want only one resource to be associated with the cross-service access, use `aws:SourceArn`. If you want to associate any resource in that AWS account with cross-service access, use `aws:SourceAccount`.

Note

The most effective way to protect against the confused deputy problem is to use the `aws:SourceArn` global condition context key with the **full ARN** of the resource. If you don't know the full ARN, or if you're specifying multiple resources, use the `aws:SourceArn` global condition context key with wildcards (*) for the unknown portions of the ARN. For example, `arn:aws:transcribe::123456789012:*`.

For an example of an assume role policy that shows how you can prevent a confused deputy issue, see [Confused deputy prevention policy \(p. 239\)](#).

Amazon Transcribe identity-based policy examples

By default, IAM users and roles don't have permission to create or modify Amazon Transcribe resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API. An IAM administrator must create IAM policies that grant users and roles permission to perform actions on the resources that they need. The administrator must then attach those policies to the IAM users or groups that require those permissions.

To learn how to create an IAM identity-based policy using these example JSON policy documents, see [Creating IAM policies](#) in the *IAM User Guide*.

Topics

- [Policy best practices \(p. 235\)](#)
- [Using the AWS Management Console \(p. 235\)](#)
- [AWS-managed \(predefined\) policies for Amazon Transcribe \(p. 236\)](#)
- [Permissions required for IAM user roles \(p. 236\)](#)
- [Permissions required for Amazon S3 encryption keys \(p. 237\)](#)
- [Allow users to view their own permissions \(p. 238\)](#)
- [AWS KMS encryption context policy \(p. 238\)](#)
- [Confused deputy prevention policy \(p. 239\)](#)
- [Viewing transcription jobs based on tags \(p. 239\)](#)

Policy best practices

Identity-based policies are very powerful. They determine whether someone can create, access, or delete Amazon Transcribe 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 using AWS managed policies** – To start using Amazon Transcribe quickly, use AWS managed policies to give your employees the permissions they need. These policies are already available in your account and are maintained and updated by AWS. For more information, see [Get started using permissions with AWS managed policies](#) in the *IAM User Guide*.
- **Grant least privilege** – When you create custom policies, grant only the permissions required to perform a task. Start with a minimum set of permissions and grant additional permissions as necessary. Doing so is more secure than starting with permissions that are too lenient and then trying to tighten them later. For more information, see [Grant least privilege](#) in the *IAM User Guide*.
- **Enable MFA for sensitive operations** – For extra security, require IAM users to use multi-factor authentication (MFA) to access sensitive resources or API operations. For more information, see [Using multi-factor authentication \(MFA\) in AWS](#) in the *IAM User Guide*.
- **Use policy conditions for extra security** – To the extent that it's practical, define the conditions under which your identity-based policies allow access to a resource. For example, you can write conditions to specify a range of allowable IP addresses that a request must come from. You can also write conditions to allow requests only within a specified date or time range, or to require the use of SSL or MFA. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.

Using the AWS Management Console

To access the Amazon Transcribe console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon Transcribe 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 (IAM 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 you're trying to perform.

To ensure that an entity (users and roles) can use the [AWS Management Console](#), attach the following AWS-managed policy to them. For more information, see [Adding Permissions to a User](#) in the *AWS Identity and Access Management User Guide*.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "transcribe:*"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}
```

AWS-managed (predefined) policies for Amazon Transcribe

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. 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 *AWS Identity and Access Management User Guide*.

The following AWS-managed policies, which you can attach to entities in your AWS account, are specific to Amazon Transcribe:

- **AmazonTranscribeReadOnly:** Grants read-only access to Amazon Transcribe resources so that you can get and list transcription jobs and custom vocabularies.
- **AmazonTranscribeFullAccess:** Grants full access to create, read, update, delete, and run all Amazon Transcribe resources. It also allows access to Amazon S3 buckets with `transcribe` in the bucket name.

Note

You can review the managed permission policies by signing in to the IAM AWS Management Console and searching by policy name. A search for "transcribe" returns both policies listed above (*AmazonTranscribeReadOnly* and *AmazonTranscribeFullAccess*).

You can also create your own custom IAM policies to allow permissions for Amazon Transcribe API actions. You can attach these custom policies to the entities that require those permissions.

Permissions required for IAM user roles

When you create an IAM user to call Amazon Transcribe, the identity must have permission to access the Amazon S3 bucket and the KMS key used to encrypt the contents of the bucket, if you provided one.

Trust policy

The IAM entity you use to make your transcription request must have a trust policy that enables Amazon Transcribe to assume that role. Use the following trust policy.

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

```
"Statement": [
  {
    "Effect": "Allow",
    "Principal": {
      "Service": [
        "transcribe.amazonaws.com"
      ]
    },
    "Action": "sts:AssumeRole"
  }
]
```

Amazon S3 input bucket policy

The following policy gives an IAM role permission to access files from the specified input bucket.

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": [
      "s3:GetObject",
      "s3:ListBucket"
    ],
    "Resource": [
      "arn:aws:s3:::DOC-EXAMPLE-INPUT-BUCKET",
      "arn:aws:s3:::DOC-EXAMPLE-INPUT-BUCKET/*"
    ]
  }
}
```

Amazon S3 output bucket policy

The following policy gives an IAM role permission to write files to the specified output bucket.

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

Permissions required for Amazon S3 encryption keys

If you're using a KMS key to encrypt an Amazon S3 bucket, include the following in the KMS key policy. This gives Amazon Transcribe access to the contents of the bucket. For more information about allowing access to KMS keys, see [Allowing external AWS accounts to access an KMS key](#) in the *AWS KMS Developer Guide*.

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

```

    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::111122223333:role/ExampleRole"
    },
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": "arn:aws:kms:us-west-2:111122223333:key/KMS-Example-KeyId"
  }
]
}

```

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": "*"
    }
  ]
}

```

AWS KMS encryption context policy

The following policy grants the IAM role “ExampleRole” permission to use the AWS KMS *Decrypt* and *Encrypt* operations for this particular KMS key. This policy works **only** for requests with at least one encryption context pair, in this case “color:indigoBlu3”. For more information on AWS KMS encryption context, see [AWS KMS encryption context \(p. 244\)](#).

```

{
  "Version": "2012-10-17",
  "Statement": [

```

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::111122223333:role/ExampleRole"
  },
  "Action": [
    "kms:Decrypt",
    "kms:DescribeKey",
    "kms:Encrypt",
    "kms:GenerateDataKey*",
    "kms:ReEncrypt*"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:EncryptionContext:color": "indigoBlu3"
    }
  }
}
```

Confused deputy prevention policy

Here's an example of an assume role policy that shows how you can use `aws:SourceArn` and `aws:SourceAccount` with Amazon Transcribe to prevent a confused deputy issue. For more information on confused deputy prevention, see [Cross-service confused deputy prevention \(p. 234\)](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "transcribe.amazonaws.com"
      },
      "Action": [
        "sts:AssumeRole",
      ],
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "111122223333"
        },
        "StringLike": {
          "aws:SourceArn": "arn:aws:transcribe:us-west-2:111122223333:*"
        }
      }
    }
  ]
}
```

Viewing transcription jobs based on tags

You can use conditions in your identity-based policy to control access to Amazon Transcribe resources based on tags. This example shows how you might create a policy that allows viewing a transcription job. However, permission is granted only if the transcription job tag `Owner` has the value of that user's user name. This policy also grants the permissions necessary to complete this action using the AWS Management Console.

You can attach this policy to the IAM users in your account. If a user named `richard-roe` attempts to view a transcription job, the transcription job must be tagged `Owner=richard-roe` or

owner=richard-roe. Otherwise they are denied access. The condition tag key `Owner` matches both `Owner` and `owner` because condition key names are not case-sensitive. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.

For more information on tagging in Amazon Transcribe, see [Tagging resources \(p. 96\)](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ListTranscriptionJobsInConsole",
      "Effect": "Allow",
      "Action": "transcribe:ListTranscriptionJobs",
      "Resource": "*"
    },
    {
      "Sid": "ViewTranscriptionJobsIfOwner",
      "Effect": "Allow",
      "Action": "transcribe:GetTranscriptionJobs",
      "Resource": "arn:aws:transcribe:*:*:transcription-job/*",
      "Condition": {
        "StringEquals": {"aws:ResourceTag/Owner": "${aws:username}"}
      }
    }
  ]
}
```

Troubleshooting Amazon Transcribe identity and access

Use the following information to diagnose and fix common issues that you might encounter when working with Amazon Transcribe and AWS Identity and Access Management (IAM).

Topics

- [I am not authorized to perform an action in Amazon Transcribe \(p. 240\)](#)
- [I am not authorized to perform iam:PassRole \(p. 241\)](#)
- [I want to view my access keys \(p. 241\)](#)
- [I'm an administrator and want to allow others to access Amazon Transcribe \(p. 241\)](#)
- [I want to allow people outside of my AWS account to access my Amazon Transcribe resources \(p. 241\)](#)

I am not authorized to perform an action in Amazon Transcribe

If you're using the AWS Management Console, and you get a message that you're not authorized to perform an action, contact your administrator for assistance. Your administrator is the person that provided you with your user name and password.

The following example error occurs when the `mateojackson` IAM user tries to use the AWS Management Console to view details about a transcription job but doesn't have `transcribe:GetTranscriptionJob` permissions.

```
User: arn:aws:iam::111122223333:user/mateojackson is not
authorized to perform: transcribe:GetTranscriptionJob on resource:
my-transcription-job
```

In this case, Mateo asks his administrator to update his policies to allow him to access the `my-transcription-job` resource using the `transcribe:GetTranscriptionJob` action.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, then you must contact your administrator for assistance. Your administrator is the person that provided you with your user name and password. Ask that person to update your policies to allow you to pass a role to Amazon Transcribe.

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 Transcribe. However, the action requires the service to have permissions 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 asks her administrator to update her policies to allow her to perform the `iam:PassRole` action.

I want to view my access keys

After you create your IAM user access keys, you can view your access key ID at any time. However, you can't view your secret access key again. If you lose your secret key, you must create a new access key pair.

Access keys consist of two parts: an access key ID (for example, `AKIAIOSFODNN7EXAMPLE`) and a secret access key (for example, `wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY`). Like a user name and password, you must use both the access key ID and secret access key together to authenticate your requests. Manage your access keys as securely as you do your user name and password.

Important

Do not provide your access keys to a third party, even to help [find your canonical user ID](#). By doing this, you might give someone permanent access to your account.

When you create an access key pair, you are prompted to save the access key ID and secret access key in a secure location. The secret access key is available only at the time you create it. If you lose your secret access key, you must add new access keys to your IAM user. You can have a maximum of two access keys. If you already have two, you must delete one key pair before creating a new one. To view instructions, see [Managing access keys](#) in the *IAM User Guide*.

I'm an administrator and want to allow others to access Amazon Transcribe

To allow others to access Amazon Transcribe, you must create an IAM entity (user or role) for the person or application that needs access. They will use the credentials for that entity to access AWS. You must then attach a policy to the entity that grants them the correct permissions in Amazon Transcribe.

To get started right away, see [Creating your first IAM delegated user and group](#) in the *IAM User Guide*.

I want to allow people outside of my AWS account to access my Amazon Transcribe 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 Transcribe supports these features, see [How Amazon Transcribe works with IAM](#) (p. 229).
- 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 [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.

Data protection in Amazon Transcribe

The AWS [shared responsibility model](#) applies to data protection in Amazon Transcribe. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. This content includes the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual user accounts with AWS Identity and Access Management (IAM). That way each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We recommend TLS 1.2 or later.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing personal data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-2](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form fields such as a **Name** field. This includes when you work with Amazon Transcribe or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Inter-network traffic privacy

An Amazon Virtual Private Cloud (Amazon VPC) endpoint for Amazon Transcribe is a logical entity within a VPC that allows connectivity only to Amazon Transcribe. Amazon VPC routes requests to Amazon Transcribe and routes responses back to the VPC. For more information, see [AWS PrivateLink concepts](#).

For information about using Amazon VPC endpoints with Amazon Transcribe see [Amazon Transcribe and interface VPC endpoints \(AWS PrivateLink\) \(p. 253\)](#).

Data encryption

Data encryption refers to protecting data while in transit and at rest. You can protect your data by using Amazon S3-managed keys or KMS keys at rest, alongside standard Transport Layer Security (TLS) while in transit.

Encryption at rest

Amazon Transcribe uses the default Amazon S3 key (SSE-S3) for server-side encryption of transcripts placed in your Amazon S3 bucket.

When you use the [StartTranscriptionJob](#) operation, you can specify your own KMS key to encrypt the output from a transcription job.

Amazon Transcribe uses an Amazon EBS volume encrypted with the default key.

Encryption in transit

Amazon Transcribe uses TLS 1.2 with AWS certificates to encrypt data in transit. This includes streaming transcriptions.

Key management

Amazon Transcribe works with KMS keys to provide enhanced encryption for your data. Amazon S3 already enables you to encrypt your input media when creating a transcription job. Integration with AWS KMS enables you to encrypt the output from a [StartTranscriptionJob](#) request.

If you don't specify a KMS key, the output of the transcription job is encrypted with the default Amazon S3 key (SSE-S3).

For more information on AWS KMS, see the [AWS Key Management Service Developer Guide](#).

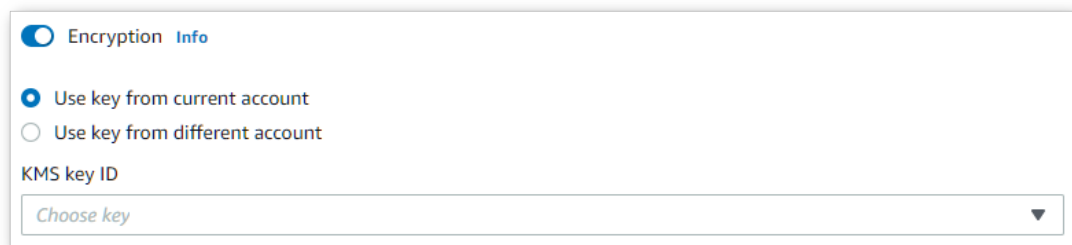
Key management using the AWS Management Console

To encrypt the output of your transcription job, you can choose between using a KMS key for the AWS account that is making the request, or you can use a KMS key from another AWS account.

If you don't specify a KMS key, the output of the transcription job is encrypted with the default Amazon S3 key (SSE-S3).

To enable output encryption:

1. Under **Output data** choose **Encryption**.



Encryption [Info](#)

Use key from current account

Use key from different account

KMS key ID

Choose key ▼

2. Choose whether the KMS key is from the AWS account you're currently using or from a different AWS account. If you want to use a key from the current AWS account, choose the key from **KMS key ID**. If

you're using a key from a different AWS account, you need to enter the key's ARN. To use a key from a different AWS account, the caller must have `kms:Encrypt` permissions for the KMS key. Refer to [Creating a key policy](#) for more information.

Key management using the API

To use output encryption with the API, you must specify your KMS key using the `OutputEncryptionKMSKeyId` parameter of the [StartCallAnalyticsJob](#), [StartMedicalTranscriptionJob](#), or [StartTranscriptionJob](#) operation.

If using a key located in the **current** AWS account, you can specify your KMS key in one of four ways:

1. Use the KMS key ID itself. For example, `1234abcd-12ab-34cd-56ef-1234567890ab`.
2. Use an alias for the KMS key ID. For example, `alias/ExampleAlias`.
3. Use the Amazon Resource Name (ARN) for the KMS key ID. For example, `arn:aws:kms:region:account-ID:key/1234abcd-12ab-34cd-56ef-1234567890ab`.
4. Use the ARN for the KMS key alias. For example, `arn:aws:kms:region:account-ID:alias/ExampleAlias`.

If using a key located in a **different** AWS account than the current AWS account, you can specify your KMS key in one of two ways:

1. Use the ARN for the KMS key ID. For example, `arn:aws:kms:region:account-ID:key/1234abcd-12ab-34cd-56ef-1234567890ab`.
2. Use the ARN for the KMS key alias. For example, `arn:aws:kms:region:account-ID:alias/ExampleAlias`.

Note that the user making the request must have permission to use the specified KMS key.

AWS KMS encryption context

AWS KMS encryption context is a map of plain text, non-secret key:value pairs. This map represents additional authenticated data, known as encryption context pairs, which provide an added layer of security for your data. Amazon Transcribe requires a symmetric encryption key to encrypt transcription output into a customer-specified Amazon S3 bucket. To learn more, see [Asymmetric keys in AWS KMS](#).

When creating your encryption context pairs, **do not** include sensitive information. Encryption context is not secret—it is visible in plain text within your CloudTrail logs (so you can use it to identify and categorize your cryptographic operations).

Your encryption context pair can include special characters, such as underscores (`_`), dashes (`-`), slashes (`/`, `\`) and colons (`:`).

Tip

It can be useful to relate the values in your encryption context pair to the data being encrypted. Although not required, we recommend you use non-sensitive metadata related to your encrypted content, such as file names, header values, or unencrypted database fields.

To use output encryption with the API, set the `KMSEncryptionContext` parameter in the [StartTranscriptionJob](#) operation. In order to provide encryption context for the output encryption operation, the `OutputEncryptionKMSKeyId` parameter must reference a symmetric KMS key ID.

You can use [AWS KMS condition keys](#) with IAM policies to control access to a symmetric encryption KMS key based on the encryption context that was used in the request for a [cryptographic operation](#). For an example encryption context policy, see [AWS KMS encryption context policy \(p. 238\)](#).

Using encryption context is optional, but recommended. For more information, see [Encryption context](#).

Opting out of using your data for service improvement

By default, Amazon Transcribe stores and uses voice inputs that it has processed to develop the service and continuously improve your experience. You can opt out of having your content used to develop and improve Amazon Transcribe by using an AWS Organizations opt-out policy. For information about how to opt out, see [AI services opt-out policies](#).

Opting out results in:

- Deletion of all transcripts stored in service-managed buckets.
- Disallowed future use of service-managed buckets.

After opting out, you must specify an output location (`OutputBucketName` or `OutputLocation`) for your transcripts when you make a [StartTranscriptionJob](#), [StartMedicalTranscriptionJob](#), or [StartCallAnalyticsJob](#) request. If you don't specify an output location, you get a `BadRequestException` error.

If you want to keep any transcripts stored in service-managed buckets, you must download them prior to opting out. To see which of your transcripts are in service-managed buckets, make a [ListTranscriptionJobs](#) request and note the location specified in `OutputLocationType`.

Monitoring Amazon Transcribe

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon Transcribe and your other AWS solutions. AWS provides the following monitoring tools to watch Amazon Transcribe, report when something is wrong, and take automatic actions when appropriate:

- **Amazon CloudWatch** monitors your AWS resources and the applications you run on AWS in real time. You can collect and track metrics, create customized dashboards, and set alarms that notify you or take actions when a specified metric reaches a threshold that you specify. For example, you can have CloudWatch track CPU usage or other metrics of your Amazon EC2 instances and automatically launch new instances when needed.
- **Amazon CloudWatch Logs** enables you to monitor, store, and access your log files from Amazon EC2 instances, CloudTrail, and other sources. CloudWatch Logs can monitor information in the log files and notify you when certain thresholds are met. You can also archive your log data in highly durable storage.
- **AWS CloudTrail** captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred.

For more information, see the [Amazon CloudWatch User Guide](#).

Amazon EventBridge is a serverless event bus service that makes it easy to connect your applications with data from a variety of sources. EventBridge delivers a stream of real-time data from your own applications, Software-as-a-Service (SaaS) applications, and AWS services and routes that data to targets such as Lambda. This enables you to monitor events that happen in services, and build event-driven architectures. For more information, see the [Amazon EventBridge User Guide](#).

Monitoring Amazon Transcribe with Amazon CloudWatch

You can monitor Amazon Transcribe using CloudWatch, which collects raw data and processes it into readable, near real-time metrics. These statistics are kept for 15 months, so that you can access historical information and gain a better perspective on how your web application or service is performing. You can also set alarms that watch for certain thresholds, and send notifications or take actions when those thresholds are met. For more information, see the [CloudWatch User Guide](#).

Using Amazon CloudWatch metrics and dimensions with Amazon Transcribe

Amazon Transcribe supports CloudWatch metrics and dimensions, which are data that can help you monitor performance; supported metrics categories include traffic, errors, data transfer, and latency associated with your transcription jobs. Supported metrics are located through CloudWatch in the **AWS/Transcribe** namespace.

Note

CloudWatch monitoring metrics are free of charge and don't count against CloudWatch service quotas.

For more information on CloudWatch metrics, see [Using Amazon CloudWatch metrics](#).

Monitoring Amazon Transcribe with AWS CloudTrail

Amazon Transcribe is integrated with AWS CloudTrail, a service that provides a record of actions taken in Amazon Transcribe by an AWS Identity and Access Management (IAM) user or role, or by an AWS service. CloudTrail captures all API calls for Amazon Transcribe, including calls from the AWS Management Console and from code calls to the Amazon Transcribe APIs, as events. By creating a trail, you can enable continuous delivery of CloudTrail events, including events for Amazon Transcribe, to an Amazon S3 bucket. If you don't create a trail, you can still view the most recent events in the CloudTrail AWS Management Console in **Event history**. Using the information collected by CloudTrail, you can see each request that is made to Amazon Transcribe, the IP address from which the request is made, who made the request, when it is made, and additional details.

To learn more about CloudTrail, refer to the [AWS CloudTrail User Guide](#).

Amazon Transcribe and CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Amazon Transcribe, that activity is recorded in a CloudTrail event along with other AWS service events in the CloudTrail **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing Events with CloudTrail Event History](#).

To get an ongoing record of events in your AWS account, including events for Amazon Transcribe, create a trail. A *trail* is a configuration that enables CloudTrail to deliver events as log files to a specified Amazon S3 bucket. 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 aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

By default, when you create a trail in the AWS Management Console, the trail applies to all AWS Regions. The trail logs events from all AWS Regions in the AWS partition and delivers the log files to the Amazon

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](#)

CloudTrail logs all Amazon Transcribe actions, which are documented in the [Reference](#). For example, calls to the [CreateVocabulary](#), [GetTranscriptionJob](#), and [StartTranscriptionJob](#) APIs generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. This information helps you determine the following:

- Whether the request is made with root or IAM user credentials
- Whether the request is made with temporary security credentials for an IAM role or federated user
- Whether the request is made by another AWS service

For more information, see the [CloudTrail userIdentity Element](#).

You can also aggregate Amazon Transcribe log files from multiple AWS Regions and multiple AWS accounts into a single Amazon S3 bucket. For more information, see [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#).

Example: Amazon Transcribe log file entries

A *trail* is a configuration that enables delivery of events as log files to a specified Amazon S3 bucket. 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 aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

Calls to the [StartTranscriptionJob](#) and [GetTranscriptionJob](#) APIs create the following entry.

```
{
  "Records": [
    {
      "eventVersion": "1.05",
      "userIdentity": {
        "type": "IAMUser",
        "principalId": "111122223333",
        "arn": "arn:aws:iam:us-west-2:111122223333:user/my-user-name",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "my-user-name"
      },
      "eventTime": "2022-03-07T15:03:45Z",
      "eventSource": "transcribe.amazonaws.com",
      "eventName": "StartTranscriptionJob",
      "awsRegion": "us-west-2",
      "sourceIPAddress": "127.0.0.1",
      "userAgent": "[ ]",
      "requestParameters": {
        "mediaFormat": "flac",
        "languageCode": "en-US",

```

```
        "transcriptionJobName": "my-first-transcription-job",
        "media": {
            "mediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-media-file.flac"
        }
    },
    "responseElements": {
        "transcriptionJob": {
            "transcriptionJobStatus": "IN_PROGRESS",
            "mediaFormat": "flac",
            "creationTime": "2022-03-07T15:03:44.229000-08:00",
            "transcriptionJobName": "my-first-transcription-job",
            "languageCode": "en-US",
            "media": {
                "mediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-media-file.flac"
            }
        }
    },
    "requestID": "47B8E8D397DCE7A6",
    "eventID": "cdc4b7ed-e171-4cef-975a-ad829d4123e8",
    "eventType": "AwsApiCall",
    "recipientAccountId": "111122223333"
},
{
    "eventVersion": "1.05",
    "userIdentity": {
        "type": "IAMUser",
        "principalId": "111122223333",
        "arn": "arn:aws:iam:us-west-2:111122223333:user/my-user-name",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "my-user-name"
    },
    "eventTime": "2022-03-07T15:07:11Z",
    "eventSource": "transcribe.amazonaws.com",
    "eventName": "GetTranscriptionJob",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "127.0.0.1",
    "userAgent": "[ ]",
    "requestParameters": {
        "transcriptionJobName": "my-first-transcription-job"
    },
    "responseElements": {
        "transcriptionJob": {
            "settings": {
            },
            "transcriptionJobStatus": "COMPLETED",
            "mediaFormat": "flac",
            "creationTime": "2022-03-07T15:03:44.229000-08:00",
            "transcriptionJobName": "my-first-transcription-job",
            "languageCode": "en-US",
            "media": {
                "mediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-media-file.flac"
            },
            "transcript": {
                "transcriptFileUri": "s3://DOC-EXAMPLE-BUCKET/my-first-
transcription-job.json"
            }
        }
    },
    "requestID": "BD8798EACDD16751",
    "eventID": "607b9532-1423-41c7-b048-ec2641693c47",
    "eventType": "AwsApiCall",
    "recipientAccountId": "111122223333"
}
]
```

```
}
```

Using Amazon EventBridge with Amazon Transcribe

With Amazon EventBridge, you can respond to state changes in your Amazon Transcribe jobs by initiating events in other AWS services. When a transcription job changes state, EventBridge automatically sends an event to an event stream. You create rules that define the events that you want to monitor in the event stream and the action that EventBridge should take when those events occur. For example, routing the event to another service (or target), which can then take an action. You could, for example, configure a rule to route an event to an AWS Lambda function when a transcription job has completed successfully.

Before using EventBridge, you should understand the following concepts:

- **Event** – An event indicates a change in the state of one of your transcription jobs. For example, when the `TranscriptionJobStatus` of a job changes from `IN_PROGRESS` to `COMPLETED`.
- **Target** – A target is another AWS service that processes an event. For example, AWS Lambda or Amazon Simple Notification Service (Amazon SNS). A target receives events in JSON format.
- **Rule** – A rule matches incoming events that you want EventBridge to watch for and routes them to a target or targets for processing. If a rule routes an event to multiple targets, all of the targets process the event in parallel. A rule can customize the JSON sent to the target.

Amazon EventBridge are emitted on a best-effort basis. For more information about creating and managing events in EventBridge, see [Amazon EventBridge events](#) in the *Amazon EventBridge User Guide*.

Defining EventBridge rules

To define EventBridge rules, use the [AWS Management Console](#). When you define a rule, use Amazon Transcribe as the service name. For an example of how to create an EventBridge rule, see [Amazon EventBridge rules](#).

The following is an example of an EventBridge rule for Amazon Transcribe that's initiated when a transcription job's status changes to `COMPLETED` or `FAILED`.

```
{
  "source": [
    "aws.transcribe"
  ],
  "detail-type": [
    "Transcribe Job State Change"
  ],
  "detail": {
    "TranscriptionJobStatus": [
      "COMPLETED",
      "FAILED"
    ]
  }
}
```

The rule contains the following fields:

- `source` – The source of the event. For Amazon Transcribe, this is always `aws.transcribe`.
- `detail-type` – An identifier for the details of the event. For Amazon Transcribe, this is always `Transcribe Job State Change`.
- `detail` – The new job status of the transcription job. In this example, the rule initiates an event when the job status changes to `COMPLETED` or `FAILED`. For a list of status values, see the `TranscriptionJobStatus` field of the [TranscriptionJob](#) API.

Amazon Transcribe events

Amazon EventBridge logs three kinds of Amazon Transcribe events: transcription job events, automatic language identification events, and Call Analytics events.

Definitions for all fields are provided below the event examples.

Transcription job events

When a job's state changes from `IN_PROGRESS` to either `COMPLETED` or `FAILED`, Amazon Transcribe generates an event. To identify the job that changed state and initiates the event in your target, use the event's `TranscriptionJobName` field. An Amazon Transcribe event contains the following information. A `FailureReason` field is added under `detail` if your transcription job status is `FAILED`.

```
{
  "version": "0",
  "id": "event ID",
  "detail-type": "Transcribe Job State Change",
  "source": "aws.transcribe",
  "account": "111122223333",
  "time": "timestamp",
  "region": "us-west-2",
  "resources": [ ],
  "detail": {
    "TranscriptionJobName": "my-first-transcription-job",
    "TranscriptionJobStatus": "COMPLETED"
  }
}
```

Language identification events

When you enable automatic language identification, Amazon Transcribe generates an event when the language identification state is either `COMPLETED` or `FAILED`. To identify the job that changed state and initiates the event in your target, use the event's `JobName` field. An Amazon Transcribe event contains the following information. A `FailureReason` field is added under `detail` if your language identification status is `FAILED`.

```
{
  "version": "0",
  "id": "event ID",
  "detail-type": "Language Identification State Change",
  "source": "aws.transcribe",
  "account": "111122223333",
  "time": "timestamp",
  "region": "us-west-2",
  "resources": [ ],
  "detail": {
    "JobType": "TranscriptionJob",
    "JobName": "job-name",
    "LanguageIdentificationStatus": "status"
  }
}
```

Call Analytics events

When a Call Analytics job state changes from `IN_PROGRESS` to either `COMPLETED` or `FAILED`, Amazon Transcribe generates an event. To identify the Call Analytics job that changed state and initiates the event in your target, use the event's `JobName` field. An Amazon Transcribe event contains the following information. A `FailureReason` field is added under `detail` if your Call Analytics job status is `FAILED`.


```
{
  "version": "0",
  "id": "event ID",
  "detail-type": "Call Analytics Job State Change",
  "source": "aws.transcribe",
  "account": "11112223333",
  "time": "timestamp",
  "region": "us-west-2",
  "resources": [ ],
  "detail": {
    "JobName": "job-name",
    "JobStatus": "status"
  }
}
```

Vocabulary events

When a custom vocabulary's state changes from `PENDING` to either `READY` or `FAILED`, Amazon Transcribe generates an event. To identify the custom vocabulary that changed state and initiates the event in your target, use the event's `VocabularyName` field. An Amazon Transcribe event contains the following information. A `FailureReason` field is added under `detail` if your custom vocabulary state is `FAILED`.

```
{
  "version": "0",
  "id": "event ID",
  "detail-type": "Vocabulary State Change",
  "source": "aws.transcribe",
  "account": "11112223333",
  "time": "timestamp",
  "region": "us-west-2",
  "resources": [ ],
  "detail": {
    "VocabularyName": "unique-vocabulary-name",
    "VocabularyState": "state"
  }
}
```

Descriptions for the preceding fields:

- `version` – The version of the event data. This value is always 0.
- `id` – A unique identifier generated by EventBridge for the event.
- `detail-type` – An identifier for the details of the event. For Amazon Transcribe, this is one of `Transcribe Job State Change`, `Language Identification State Change`, `Call Analytics Job State Change`, or `Vocabulary State Change`.
- `source` – The source of the event. For Amazon Transcribe this is always `aws.transcribe`.
- `account` – The AWS account ID of the account that generated the API call.
- `timestamp` – The date and time the event is delivered.
- `region` – The AWS Region in which the request is made.
- `resources` – The resources used by the API call. For Amazon Transcribe, this field is always empty.
- `detail` – Details about the event. The fields within the `detail` tag may vary depending on the type of event, but may contain the following fields:
 - `JobType` – For transcription jobs, this value must be `TranscriptionJob`.
 - `JobName` – The unique name for your transcription job.
 - `JobStatus` – The status of your Call Analytics transcription job. It can be either `COMPLETED` or `FAILED`.

- `TranscriptionJobName` – The unique name you chose for your transcription job.
- `TranscriptionJobStatus` – The status of the transcription job. For a list of status values, see the `TranscriptionJobStatus` field of the [TranscriptionJob](#) data type.
- `LanguageIdentificationStatus` – The status of language identification in a transcription job. It can be either `COMPLETED` or `FAILED`.
- `VocabularyName` – The unique name for your custom vocabulary.
- `VocabularyState` – The processing state of your custom vocabulary. For a list of status values, see the `VocabularyState` field of the [VocabularyInfo](#) data type.
- `FailureReason` – This field is present if the state or status changes to `FAILED`, and describes the reason for the `FAILED` state or status.

Compliance Validation for Amazon Transcribe

Third-party auditors assess the security and compliance of AWS services as part of multiple AWS compliance programs, such as SOC, PCI, FedRAMP, and HIPAA.

To learn whether or other AWS services are within the scope of specific compliance programs, see [AWS Services in Scope by Compliance Program](#). For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying baseline environments on AWS that are security and compliance focused.
- [Architecting for HIPAA Security and Compliance on Amazon Web Services](#) – This whitepaper describes how companies can use AWS to create HIPAA-eligible applications.

Note

Not all AWS services are HIPAA eligible. For more information, see the [HIPAA Eligible Services Reference](#).

- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [Evaluating Resources with Rules](#) in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub](#) – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.
- [AWS Audit Manager](#) – This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

Resilience in Amazon Transcribe

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 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](#).

Infrastructure Security in Amazon Transcribe

As a managed service, Amazon Transcribe is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

You use AWS published API calls to access Amazon Transcribe through the network. Clients must support Transport Layer Security (TLS) 1.0 or later. 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.

Vulnerability analysis and management in Amazon Transcribe

Configuration and IT controls are a shared responsibility between AWS and you, our customer. For more information, see the AWS [shared responsibility model](#).

Amazon Transcribe and interface VPC endpoints (AWS PrivateLink)

You can establish a private connection between your VPC and Amazon Transcribe by creating an *interface VPC endpoint*. Interface endpoints are powered by [AWS PrivateLink](#), a technology that enables you to privately access Amazon Transcribe APIs without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with Amazon Transcribe APIs. Traffic between your VPC and Amazon Transcribe does not leave the Amazon network.

Each interface endpoint is represented by one or more [Elastic Network Interfaces](#) in your subnets.

For more information, see [Interface VPC endpoints \(AWS PrivateLink\)](#) in the *Amazon VPC User Guide*.

Considerations for Amazon Transcribe VPC endpoints

Before you set up an interface VPC endpoint for Amazon Transcribe, ensure that you review [interface endpoint properties and limitations](#) in the *Amazon VPC User Guide*.

Amazon Transcribe supports making calls to all of its API actions from your VPC.

Creating an interface VPC endpoint for Amazon Transcribe

You can create a VPC endpoint for the Amazon Transcribe service using either the Amazon VPC AWS Management Console or AWS CLI. For more information, see [Creating an interface endpoint](#) in the *Amazon VPC User Guide*.

For batch transcriptions in Amazon Transcribe, create a VPC endpoint using the following service name:

- `com.amazonaws.us-west-2.transcribe`

For streaming transcriptions in Amazon Transcribe, create a VPC endpoint using the following service name:

- `com.amazonaws.us-west-2.transcribestreaming`

If you enable private DNS for the endpoint, you can make API requests to Amazon Transcribe using its default DNS name for the AWS Region, for example, `transcribestreaming.us-east-2.amazonaws.com`.

For more information, see [Accessing a service through an interface endpoint](#) in the *Amazon VPC User Guide*.

Creating a VPC endpoint policy for Amazon Transcribe

You can attach an endpoint policy to your VPC endpoint that controls access to either the streaming service or the batch transcription service of Amazon Transcribe. The policy specifies the following information:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which actions can be performed.

For more information, see [Controlling access to services with VPC endpoints](#) in the *Amazon VPC User Guide*.

Example: VPC endpoint policy for Amazon Transcribe batch transcription actions

The following is an example of an endpoint policy for a batch transcription in Amazon Transcribe. When attached to an endpoint, this policy grants access to the listed Amazon Transcribe actions for all principals on all resources.

```
{
  "Statement": [
    {
      "Principal": "*",
      "Effect": "Allow",
      "Action": [
        "transcribe:StartTranscriptionJob",
        "transcribe:ListTranscriptionJobs"
      ],
      "Resource": "*"
    }
  ]
}
```

Example: VPC endpoint policy for Amazon Transcribe streaming transcription actions

The following is an example of an endpoint policy for a streaming transcription in Amazon Transcribe. When attached to an endpoint, this policy grants access to the listed Amazon Transcribe actions for all principals on all resources.

```
{
  "Statement": [
```

```
{
  "Principal": "*",
  "Effect": "Allow",
  "Action": [
    "transcribe:StartStreamTranscription",
    "transcribe:StartStreamTranscriptionWebsocket"
  ],
  "Resource": "*"
}
]
```

Security Best Practices for Amazon Transcribe

The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, use them as helpful considerations rather than prescriptions.

- Use data encryption, such as AWS KMS encryption context

AWS KMS encryption context is a map of plain text, non-secret key:value pairs. This map represents additional authenticated data, known as encryption context pairs, which provide an added layer of security for your data.

For more information, refer to [AWS KMS encryption context \(p. 244\)](#).

- Use IAM roles for applications and AWS services that require Amazon Transcribe access

Use an IAM role to manage temporary credentials for applications or services that need to access Amazon Transcribe. When you use a role, you don't have to distribute long-term credentials, such as passwords or access keys, to an Amazon EC2 instance or AWS service. IAM roles can supply temporary permissions that applications can use when they make requests to AWS resources.

For more information, refer to [IAM roles](#) and [Common scenarios for roles: Users, applications, and services](#).

- Use tag-based access control

You can use tags to control access within your AWS accounts. In Amazon Transcribe, tags can be added to: transcription jobs, custom vocabularies, custom vocabulary filters, and custom language models.

For more information, refer to [Tag-based access control \(p. 96\)](#).

- Use AWS monitoring tools

Monitoring is an important part of maintaining the reliability, security, availability, and performance of Amazon Transcribe and your AWS solutions. You can monitor Amazon Transcribe using CloudTrail.

For more information, refer to [Monitoring Amazon Transcribe with AWS CloudTrail \(p. 246\)](#).

- Enable AWS Config

AWS Config enables you to assess, audit, and evaluate the configurations of your AWS resources. Using AWS Config, you can review changes in configurations and relationships between AWS resources, investigate detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This can help you simplify compliance auditing, security analysis, change management, and operational troubleshooting.

For more information, refer to [What Is AWS Config?](#)

Amazon Transcribe Medical

Amazon Transcribe Medical is an automatic speech recognition (ASR) service designed for medical professionals who wish to transcribe medical-related speech, such as physician-dictated notes, drug safety monitoring, telemedicine appointments, or physician-patient conversations. Amazon Transcribe Medical is available through either real-time streaming (via microphone) or transcription of an uploaded file (batch).

Important

Amazon Transcribe Medical is not a substitute for professional medical advice, diagnosis, or treatment. Identify the right confidence threshold for your use case, and use high confidence thresholds in situations that require high accuracy. For certain use cases, results should be reviewed and verified by appropriately trained human reviewers. Amazon Transcribe Medical transcriptions should only be used in patient care scenarios after review for accuracy and sound medical judgment by trained medical professionals.

Amazon Transcribe Medical operates under a shared responsibility model, whereby AWS is responsible for protecting the infrastructure that runs Amazon Transcribe Medical and you are responsible for managing your data. For more information, see [Shared Responsibility Model](#).

Amazon Transcribe Medical is available in US English (en-US).

For analysis of your transcripts, you can use other AWS services, such as [Amazon Comprehend Medical](#).

Supported specialties

Specialty	Sub-specialty	Audio input
Cardiology	none	streaming only
Neurology	none	streaming only
Oncology	none	streaming only
Primary Care	Family Medicine	batch, streaming
Primary Care	Internal Medicine	batch, streaming
Primary Care	Obstetrics and Gynecology (OB-GYN)	batch, streaming
Primary Care	Pediatrics	batch, streaming
Radiology	none	streaming only
Urology	none	streaming only

How Amazon Transcribe Medical works

Amazon Transcribe Medical enables you to transcribe individual audio content with medical information into text.

Transcription methods can be separated into two main categories:

- Batch transcription jobs: Transcribing media that have been uploaded into an Amazon S3 bucket.
- Streaming transcriptions: Transcribing media streams in real time.

Batch transcription jobs can be created with the [AWS CLI \(p. 63\)](#), [AWS Management Console \(p. 55\)](#), and various [AWS SDKs \(p. 66\)](#).

Streaming transcriptions can be created with the [AWS Management Console \(p. 55\)](#), [HTTP/2 \(p. 82\)](#), [WebSockets \(p. 86\)](#), and various [AWS SDKs \(p. 66\)](#).

For more information on medical transcriptions with streaming audio, see [Streaming transcription overview \(p. 258\)](#) and [Establish a bi-directional connection using the WebSocket protocol \(p. 283\)](#). For more information on batch transcriptions, see [Batch transcription overview \(p. 258\)](#).

Speech input

Amazon Transcribe Medical can transcribe speech as either an audio file or a real-time stream. Your input audio must use the encodings and formats described in the following sections.

Topics

- [Containers and formats for batch transcription \(p. 257\)](#)
- [Audio containers and formats for streaming transcription \(p. 257\)](#)

Containers and formats for batch transcription

When you transcribe audio using the [StartMedicalTranscriptionJob](#) API or the AWS Management Console, make sure that the file is:

- In FLAC, MP3, MP4, Ogg, WebM, AMR, or WAV file format
- Less than 4 hours in length and less than 2 GB in size
- Encoded at a sample rate of 16,000 Hz or higher

Note

For AMR, Amazon Transcribe Medical supports both Adaptive Multi-Rate Wideband (AMR-WB) and Adaptive Multi-Rate Narrowband (AMR-NB) codecs.

For the Ogg and WebM file formats, Amazon Transcribe Medical supports the Opus codec.

For best results:

- Use a lossless format. You can choose either FLAC, or WAV with PCM 16-bit encoding.

Audio containers and formats for streaming transcription

When you transcribe a real-time stream using the [StartMedicalStreamTranscription](#) API or a WebSocket request, make sure that your stream is encoded in:

- PCM signed 16-bit little-endian
- FLAC
- OPUS encoded audio in the Ogg container

Your stream must use a sample rate of 16,000 Hz or higher.

For best results:

- Use a lossless format, such as FLAC or PCM encoding.

For more information on using a WebSocket request to transcribe your streaming audio, see [Establish a bi-directional connection using the WebSocket protocol \(p. 283\)](#).

Streaming transcription overview

Streaming transcription takes a stream of your audio data and transcribes it in real time. It uses a bidirectional WebSocket connection so that the results of the transcription are returned to your application while you send more audio to Amazon Transcribe Medical. You can also use it when you have an audio file that you want to process as it is transcribed.

Streaming transcription is available in US English (en-US). It can produce transcriptions of accented English, spoken by non-native speakers. Streaming audio transcription comes with the following features:

- Transcribe 16 kHz audio in real time.
- Transcribe up to 4 hours of audio streams.
- Word-level timestamp in transcripts.
- Word-level confidence in transcripts.
- [Number normalization \(p. 259\)](#).
- Punctuation and true casing in transcripts.
- Supports both dictation and conversation speech types.

For more information, see [Establish a bi-directional connection using the WebSocket protocol \(p. 283\)](#).

Batch transcription overview

Amazon Transcribe Medical batch transcription is available in US English. It has the ability to transcribe accented English from non-native speakers. It supports the transcription of individual audio files. You start a transcription job with either the AWS Management Console or by direct API call.

You interact with Amazon Transcribe Medical using four main API resources. To start a medical transcription job, use the [StartMedicalTranscriptionJob](#) API. To retrieve information on a medical transcription job, use [GetMedicalTranscriptionJob](#). You list medical transcription jobs with [ListMedicalTranscriptionJobs](#). You delete a medical transcription job with [DeleteMedicalTranscriptionJob](#).

To transcribe an audio file, you use a transcription job. You store the file as an object in an Amazon S3 bucket. The input file must:

- Be in FLAC, MP3, MP4, or WAV file format.
- Use 16-bit Linear PCM encoding.
- Be less than 4 hours in length and less than 2 GB in size
- Use a sample rate of 16,000 Hz or higher

For best results:

- Use a lossless format, such as FLAC or WAV.

When creating a medical transcription job, you specify the language, the medical specialty, and the audio type of the source file. You input US English (en-US) as the language and `PRIMARYCARE` as the medical specialty. Entering primary care as the value enables you to generate transcriptions from source audio in the following medical specialties:

- Family Medicine
- Internal Medicine
- Obstetrics and Gynecology (OB-GYN)
- Pediatrics

You have the choice between dictation and conversation for your audio type. Choose dictation for audio files where the physician is giving a report about a patient visit or procedure. Choose conversation for audio files that involve a conversation between a physician and a patient or a conversation between physicians.

To store the output of your transcription job, select an Amazon S3 bucket that you've already created. For more information on Amazon S3 buckets see [Getting Started with Amazon Simple Storage Service](#).

The following are the minimum number of request parameters to enter in the sample JSON.

```
{
  "MedicalTranscriptionJobName": "my-first-transcription-job",
  "LanguageCode": "en-US",
  "Media": {
    "MediaFileUri": "s3://path to your audio file"
  },
  "OutputBucketName": "your output bucket name",
  "Specialty": "PRIMARYCARE",
  "Type": "CONVERSATION"
}
```

Amazon Transcribe Medical enables you to generate alternative transcriptions. For more information, see [Generating alternative transcriptions \(p. 327\)](#).

You can also identify different speakers or channels in your audio. For more information, see [Identifying speakers and labeling their speech \(p. 296\)](#) and [Transcribing multi-channel audio \(p. 302\)](#).

Transcribing numbers

Amazon Transcribe Medical transcribes digits as numbers instead of words. For example, the spoken number "one thousand two hundred forty-two" is transcribed as 1242.

Numbers are transcribed according to the following rules.

Rule	Description
Convert cardinal numbers greater than 10 to numbers.	<ul style="list-style-type: none">• "Fifty five" > 55• "a hundred" > 100• "One thousand and thirty one" > 1031• "One hundred twenty-three million four hundred fifty six thousand seven hundred eight nine" > 123,456,789
Convert cardinal numbers followed by "million" or "billion" to numbers followed by a word when "million" or "billion" is not followed by a number.	<ul style="list-style-type: none">• "one hundred million" > 100 million• "one billion" > 1 billion• "two point three million" > 2.3 million
Convert ordinal numbers greater than 10 to numbers.	<ul style="list-style-type: none">• "Forty third" > 43rd

Rule	Description
	<ul style="list-style-type: none"> "twenty sixth avenue" > 26th avenue
Convert fractions to their numeric format.	<ul style="list-style-type: none"> "a quarter" > 1/4 "three sixteenths" > 3/16 "a half" > 1/2 "a hundredth" > 1/100
Convert numbers less than 10 to digits if there are more than one in a row.	<ul style="list-style-type: none"> "three four five" > 345 "My phone number is four two five five five five one two one two" > 4255551212
Decimals are indicated by "dot" or "point."	<ul style="list-style-type: none"> "three hundred and three dot five" > 303.5 "three point twenty three" > 3.23 "zero point four" > 0.4 "point three" > 0.3
Convert the word "percent" after a number to a percent sign.	<ul style="list-style-type: none"> "twenty three percent" > 23% "twenty three point four five percent" > 23.45%
Convert the words "dollar," "US dollar," "Australian dollar," "AUD," or "USD" after a number to a dollar symbol before the number.	<ul style="list-style-type: none"> "one dollar and fifteen cents" > \$1.15 "twenty three USD" > \$23 "twenty three Australian dollars" > \$23
Convert the words "pounds," or "milligrams" to "lbs" or "mg".	<ul style="list-style-type: none"> "twenty three pounds" > 23 lbs "forty-five milligrams" > 45 mg
Convert the words "rupees," "Indian rupees," or "INR" after a number to rupee sign (₹) before the number.	<ul style="list-style-type: none"> "twenty three rupees" > ₹23 "fifty rupees thirty paise" > ₹50.30
Convert times to numbers.	<ul style="list-style-type: none"> "seven a m eastern standard time" > 7 a.m. eastern standard time "twelve thirty p m" > 12:30 p.m.
Combine years expressed as two digits into four. Only valid for the 20th, 21st, and 22nd centuries.	<ul style="list-style-type: none"> "nineteen sixty two" > 1962 "the year is twenty twelve" > the year is 2012 "twenty nineteen" > 2019 "twenty one thirty" > 2130
Convert dates to numbers.	<ul style="list-style-type: none"> "May fifth twenty twelve" > May 5th 2012 "May five twenty twelve" > May 5 2012 "five May twenty twelve" > 5 May 2012
Separate spans of numbers by the word "to."	<ul style="list-style-type: none"> "twenty three to thirty seven" > 23 to 37

Transcribing medical terms and measurements

Amazon Transcribe Medical can transcribe medical terms and measurements. Amazon Transcribe Medical outputs abbreviations for spoken terms. For example, "blood pressure" is transcribed as BP. You can find a list of conventions that Amazon Transcribe Medical uses for medical terms and measurements in the table on this page. The *Spoken Term* column refers to the term spoken in the source audio. The *Output* column refers to the abbreviation you see in your transcription results.

You can see how the terms spoken in source audio correspond to the transcription output here.

Term spoken in source audio	Abbreviation used in output	Example output
Centigrade	C	The patient's temperature is 37.4 C.
Celsius	C	The patient's temperature is 37.4 C.
Fahrenheit	F	The patient's temperature is 101 F.
grams	g	A mass of 100 g was extracted from the patient.
meters	m	The patient is 1.8 m tall.
feet	ft	The patient is 6 ft tall.
kilos	kg	The patient weighs 80 kg.
kilograms	kg	The patient weighs 80 kg.
c c	cc	Patient received 100 cc of saline solution.
cubic centimeter	cc	Patient received 100 cc of saline solution.
milliliter	mL	Patient excreted 100 mL urine.
blood pressure	BP	Patient BP was elevated.
b p	BP	Patient BP was elevated.
X over Y	X/Y	Patient BP was 120/80.
beats per min	BPM	Patient had atrial fibrillation with heart rate of 160 BPM.
beats per minute	BPM	Patient had atrial fibrillation with heart rate of 160 BPM.
O 2	O2	Patient O2 saturation was 98%.
CO2	CO2	Patient required respiratory support for elevated CO2.
post operation	POSTOP	Patient came for POSTOP evaluation.
post op	POSTOP	Patient came for POSTOP evaluation.
cat scan	CT Scan	Patient indication of cerebral hemorrhage required use of CT Scan.
Pulse 80	P 80	Patient vitals were P 80, R 17,...

Term spoken in source audio	Abbreviation used in output	Example output
Respiration 17	R 17	Patient vitals were P 80, R 17,...
in and out	I/O	Patient was I/O sinus rhythm
L five	L5	Lumbar puncture was performed between L4 and L5

Getting started with Amazon Transcribe Medical

To get started using Amazon Transcribe Medical, set up an AWS account and create an AWS Identity and Access Management (IAM) user.

Topics

- [Set up an AWS account and create an administrator user \(p. 262\)](#)
- [Getting started with the AWS Management Console \(streaming\) \(p. 263\)](#)
- [Getting started with batch transcription \(p. 264\)](#)

Set up an AWS account and create an administrator user

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

1. [Sign up for Amazon Web Services \(p. 262\)](#)
2. [Create an IAM user \(p. 263\)](#)

Sign up for Amazon Web Services

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

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

If you already have an AWS account, skip to the next section.

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 and entering a verification code on the phone keypad.

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

Create an IAM user

Services in AWS, such as Amazon Transcribe Medical, require that you provide credentials when you access them. This allows the service to determine whether you have permissions to access the service's resources.

We strongly recommend that you access AWS using AWS Identity and Access Management (IAM), not the credentials for your AWS account. To use IAM to access AWS, create an IAM user, add the user to an IAM group with administrative permissions, and then grant administrative permissions to the IAM user. You can then access AWS using a special URL and the IAM user's credentials.

The Getting Started exercises in this guide assume that you have a user with administrator privileges, `adminuser`.

To create an administrator user and sign in to the AWS Management Console

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

To get started with medical transcription using the AWS Management Console, see [Getting Started \(AWS Management Console\)](#) (p. 263).

Getting started with the AWS Management Console (streaming)

Get started with Amazon Transcribe Medical streaming transcription by using the AWS Management Console to transcribe up to 15 minutes of medical speech into text. You use the microphone on your computer for the audio source.

To start a streaming medical transcription job

1. Open the [AWS Management Console](#).
2. From the left menu, look under **Amazon Transcribe Medical** and choose **Real-time transcription**.
3. Choose **Start streaming** and then speak into your microphone. Amazon Transcribe Medical will transcribe the speech and display the results on the AWS Management Console.

Next step

To learn more about WebSocket streaming, go to [WebSocket Streaming](#) (p. 283).

Getting started with batch transcription

Get started with Amazon Transcribe Medical batch transcription in the AWS Management Console to transcribe your audio file. To start a transcription job, your audio file must be stored in Amazon S3 bucket and its file format must be either FLAC, WAV, MP3, or MP4.

To start a batch medical transcription job

1. Open the [AWS Management Console](#).
2. From the left menu, look under **Amazon Transcribe Medical** and choose **Transcription jobs**.
3. Choose **Create job**
4. Specify the details of your job. Choose the **Info** links if you need help
5. When you are finished specifying the details of your job, choose **Create** to start a medical transcription job.

[Batch transcription overview \(p. 258\)](#)

Streaming transcription

Amazon Transcribe Medical streaming transcription enables you to send an audio stream and receive a stream of text in real time. The API makes it easy for developers to add real-time speech-to-text capability to their applications.

The following table shows which languages are available for streaming transcription and how you can access them.

Language	Sample rate	Available in
US English (en-US)	16 kHz, 8 kHz	AWS Management Console , StartMedicalStreamTranscription API , and WebSocket request

If you are using HTTP/2, we provide an HTTP/2 streaming client that handles retrying the connection when there are transient problems on the network. You can use this client as a starting point for your own applications. To use Amazon Transcribe Medical streaming with the WebSocket protocol, you can create your own client.

Streaming transcription takes a stream of your audio data and transcribes it in real time. The transcription is returned to your application in a stream of transcription events.

Amazon Transcribe Medical breaks your incoming audio stream based on natural speech segments, such as a change in speaker or a pause in the audio. The transcription is returned progressively to your application, with each response containing more transcribed speech until the entire segment is transcribed.

In the following example, each line is a partial result transcription output of an audio segment being streamed.

```
The  
The PE.
```

```
The pain.  
The patient.  
The patient was  
The patient was in  
The patient was entered.  
The patient was entered, and, uh  
The patient was I/O.  
The patient was I/O of  
The patient was I/O of some  
The patient was I/O sign.  
The patient was I/O of Sinus.  
The patient was I/O of Sinus rhyth.  
The patient was I/O of Sinus rhythm.  
The patient was I/O of Sinus rhythm with
```

Each [Result](#) object in the response contains a field called `IsPartial` that indicates whether the response is a partial response containing the transcription results so far or if it is a complete transcription of the audio segment.

Each [Result](#) object also contains the start time and end time of the term from the audio stream so that you can, for example, synchronize the transcription with the video.

The following example is a partial transcription response.

```
{  
  "Transcript": {  
    "Results": [  
      {  
        "Alternatives": [  
          {  
            "Items": [  
              {  
                "Content": "The",  
                "EndTime": 1.07,  
                "StartTime": 1.04,  
                "Type": "pronunciation",  
                "VocabularyFilterMatch": false  
              },  
              {  
                "Content": "patient",  
                "EndTime": 1.5,  
                "StartTime": 1.08,  
                "Type": "pronunciation",  
                "VocabularyFilterMatch": false  
              },  
              {  
                "Content": "was",  
                "EndTime": 1.61,  
                "StartTime": 1.51,  
                "Type": "pronunciation",  
                "VocabularyFilterMatch": false  
              },  
              {  
                "Content": "I/O",  
                "EndTime": 2.25,  
                "StartTime": 2.06,  
                "Type": "pronunciation",  
                "VocabularyFilterMatch": false  
              },  
              {  
                "Content": "of",  
                "EndTime": 2.34,  
                "StartTime": 2.26,
```

```

        "Type": "pronunciation",
        "VocabularyFilterMatch": false
    },
    {
        "Content": "Sinus",
        "EndTime": 2.71,
        "StartTime": 2.35,
        "Type": "pronunciation",
        "VocabularyFilterMatch": false
    },
    {
        "Content": "rhythm",
        "EndTime": 3.07,
        "StartTime": 2.72,
        "Type": "pronunciation",
        "VocabularyFilterMatch": false
    },
    {
        "Content": "with",
        "EndTime": 3.68,
        "StartTime": 3.49,
        "Type": "pronunciation",
        "VocabularyFilterMatch": false
    }
    ],
    "Transcript": "The patient was I/O of Sinus rhythm with"
}
],
"EndTime": 3.75,
"IsPartial": true,
"ResultId": "93b1df2b-8702-4c91-892a-ace3b65a6477",
"StartTime": 1.04
}
]
}
}

```

The following example shows the transcription results for a fully transcribed speech segment.

```

{
  "Transcript": {
    "Results": [
      {
        "Alternatives": [
          {
            "Items": [
              {
                "Confidence": 0.99,
                "Content": "The",
                "EndTime": 1.12,
                "StartTime": 1.04,
                "Type": "pronunciation",
                "VocabularyFilterMatch": false
              },
              {
                "Confidence": 0.99,
                "Content": "patient",
                "EndTime": 1.53,
                "StartTime": 1.13,
                "Type": "pronunciation",
                "VocabularyFilterMatch": false
              },
              {
                "Confidence": 1,
                "Content": "was",

```



```
        "EndTime": 1.73,  
        "StartTime": 1.54,  
        "Type": "pronunciation",  
        "VocabularyFilterMatch": false  
    },  
    {  
        "Confidence": 1,  
        "Content": "I/O",  
        "EndTime": 2.3,  
        "StartTime": 2.12,  
        "Type": "pronunciation",  
        "VocabularyFilterMatch": false  
    },  
    {  
        "Confidence": 1,  
        "Content": "of",  
        "EndTime": 2.39,  
        "StartTime": 2.31,  
        "Type": "pronunciation",  
        "VocabularyFilterMatch": false  
    },  
    {  
        "Confidence": 1,  
        "Content": "Sinus",  
        "EndTime": 2.82,  
        "StartTime": 2.4,  
        "Type": "pronunciation",  
        "VocabularyFilterMatch": false  
    },  
    {  
        "Confidence": 0.99,  
        "Content": "rhythm",  
        "EndTime": 3.32,  
        "StartTime": 2.83,  
        "Type": "pronunciation",  
        "VocabularyFilterMatch": false  
    },  
    {  
        "Confidence": 0.99,  
        "Content": "with",  
        "EndTime": 3.72,  
        "StartTime": 3.49,  
        "Type": "pronunciation",  
        "VocabularyFilterMatch": false  
    },  
    {  
        "Confidence": 0.99,  
        "Content": "a",  
        "EndTime": 3.78,  
        "StartTime": 3.73,  
        "Type": "pronunciation",  
        "VocabularyFilterMatch": false  
    },  
    {  
        "Confidence": 1,  
        "Content": "heart",  
        "EndTime": 4.02,  
        "StartTime": 3.79,  
        "Type": "pronunciation",  
        "VocabularyFilterMatch": false  
    },  
    {  
        "Confidence": 1,  
        "Content": "rate",  
        "EndTime": 4.19,  
        "StartTime": 4.03,
```

```

        "Type": "pronunciation",
        "VocabularyFilterMatch": false
    },
    {
        "Confidence": 0.99,
        "Content": "of",
        "EndTime": 4.26,
        "StartTime": 4.2,
        "Type": "pronunciation",
        "VocabularyFilterMatch": false
    },
    {
        "Confidence": 1,
        "Content": "75",
        "EndTime": 4.81,
        "StartTime": 4.27,
        "Type": "pronunciation",
        "VocabularyFilterMatch": false
    },
    {
        "Confidence": 0.97,
        "Content": "bpm",
        "EndTime": 5.47,
        "StartTime": 4.82,
        "Type": "pronunciation",
        "VocabularyFilterMatch": false
    },
    {
        "Content": ".",
        "EndTime": 5.47,
        "StartTime": 5.47,
        "Type": "punctuation",
        "VocabularyFilterMatch": false
    }
    ],
    "Transcript": "The patient was I/O of Sinus rhythm with a heart
rate of 75 bpm."
}
    ],
    "EndTime": 5.53,
    "IsPartial": false,
    "ResultId": "93b1df2b-8702-4c91-892a-ace3b65a6477",
    "StartTime": 1.04
}
]
}
}

```

Each word, phrase or punctuation mark in the transcription output is an *item*. Each word or phrase has a *confidence* score. The confidence score is a value between 0 and 1 that indicates how confident Amazon Transcribe Medical is that it correctly transcribed the item. A confidence score with a larger value indicates that Amazon Transcribe Medical is more confident that it transcribed the item correctly.

The preceding example shows "I/O" in the transcription output, which is an abbreviation for "in and out". For more information on how Amazon Transcribe Medical uses abbreviations in the transcription output, see [Transcribing medical terms and measurements \(p. 260\)](#).

Topics

- [Event stream encoding \(p. 269\)](#)
- [Using Amazon Transcribe Medical streaming with HTTP/2 \(p. 270\)](#)
- [Establish a bi-directional connection using the WebSocket protocol \(p. 283\)](#)

Event stream encoding

Event stream encoding provides bidirectional communication using messages between a client and a server. Data frames sent to the Amazon Transcribe Medical streaming service are encoded in this format. The response from Amazon Transcribe Medical also uses this encoding.

Each message consists of two sections: the prelude and the data. The prelude consists of:

1. The total byte length of the message
2. The combined byte length of all of the headers

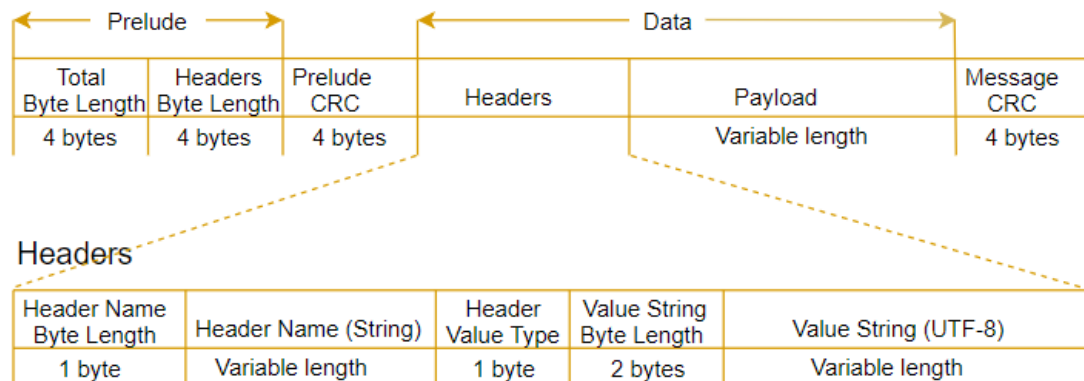
The data section consists of:

1. The headers
2. A payload

Each section ends with a 4-byte big-endian integer CRC checksum. The message CRC checksum is for both the prelude section and the data section. Amazon Transcribe Medical uses CRC32 (often referred to as GZIP CRC32) to calculate both CRCs. For more information about CRC32, see [GZIP file format specification version 4.3](#).

Total message overhead, including the prelude and both checksums, is 16 bytes.

The following diagram shows the components that make up a message and a header. There are multiple headers per message.



Each message contains the following components:

- **Prelude:** Always a fixed size of 8 bytes, two fields of 4 bytes each.
 - *First 4 bytes:* The total byte-length. This is the big-endian integer byte-length of the entire message, including the 4-byte length field itself.
 - *Second 4 bytes:* The headers byte-length. This is the big-endian integer byte-length of the headers portion of the message, excluding the headers length field itself.
- **Prelude CRC:** The 4-byte CRC checksum for the prelude portion of the message, excluding the CRC itself. The prelude has a separate CRC from the message CRC to ensure that Amazon Transcribe Medical can detect corrupted byte-length information immediately without causing errors such as buffer overruns.
- **Headers:** Metadata annotating the message, such as the message type, content type, and so on. Messages have multiple headers. Headers are key-value pairs where the key is a UTF-8 string. Headers can appear in any order in the headers portion of the message and any given header can appear only once. For the required header types, see the following sections.

- **Payload:** The audio content to be transcribed.
- **Message CRC:** The 4-byte CRC checksum from the start of the message to the start of the checksum. That is, everything in the message except the CRC itself.

Each header contains the following components. There are multiple headers per frame.

- **Header name byte-length:** The byte-length of the header name.
- **Header name:** The name of the header indicating the header type. For valid values, see the following frame descriptions.
- **Header value type:** An enumeration indicating the header value.

The following shows the possible values for the header and what they indicate.

- 0 – TRUE
- 1 – FALSE
- 2 – BYTE
- 3 – SHORT
- 4 – INTEGER
- 5 – LONG
- 6 – BYTE ARRAY
- 7 – STRING
- 8 – TIMESTAMP
- 9 – UUID
- **Value string byte length:** The byte-length of the header value string.
- **Header value:** The value of the header string. Valid values for this field depend on the type of header. For valid values, see the following frame descriptions.

Using Amazon Transcribe Medical streaming with HTTP/2

Amazon Transcribe Medical uses a format called *event stream encoding* for streaming-med transcription. This format encoded binary data with header information that describes the contents of each event. For more information, see [Event stream encoding \(p. 80\)](#). You can use this information for applications that call the Amazon Transcribe Medical endpoint without using an SDK.

When Amazon Transcribe Medical uses the [HTTP/2 protocol](#) for streaming medical transcriptions, the key components for a streaming medical request are:

- A header frame. This contains the HTTP/2 headers for the request, and a signature in the `authorization` header that Amazon Transcribe Medical uses as a seed signature to sign the following data frames.
- One or more message frames in event stream encoding. The frame contains metadata and the raw audio bytes.
- An end frame. This is a signed message in event stream encoding with an empty body.

Streaming request

To make a streaming request, you use the [StartMedicalStreamTranscription](#) API.

Header frame

The header frame is the authorization frame for the streaming transcription. Amazon Transcribe Medical uses the value of the `authorization` header as the seed for generating a chain of authorization headers for the data frames in the request.

Required headers

The header frame of a request to Amazon Transcribe Medical requires the following HTTP/2 headers.

```
POST /medical-stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
authorization: Generated value
x-amz-target: com.amazonaws.transcribe.Transcribe.StartMedicalStreamTranscription
x-amz-content-sha256: STREAMING-MEDAWS4-HMAC-SHA256-EVENTS
x-amz-date: 20220208T235959Z
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
Content-type: application/vnd.amazon.eventstream
transfer-encoding: chunked
```

In the request, use the following values for the `host`, `authorization`, and `x-amz-date` headers:

- **host:** Use the AWS Region where you are calling Amazon Transcribe Medical. For a list of valid AWS Regions, see [AWS Regions and Endpoints](#) in the *AWS General Reference*.
- **authorization:** This is a generated field. To learn more about creating a signature, see [Signing AWS requests with Signature Version 4](#).
- **x-amz-date:** The date and time the signature is created. The format is YYYYMMDDTHHMMSSZ, where YYYY=year, MM=month, DD=day, HH=hour, MM=minute, SS=seconds, and 'T' and 'Z' are fixed characters. For more information, refer to [Handling Dates in Signature Version 4](#).

For more information about the headers specific to Amazon Transcribe Medical, see the [StartMedicalStreamTranscription](#) API.

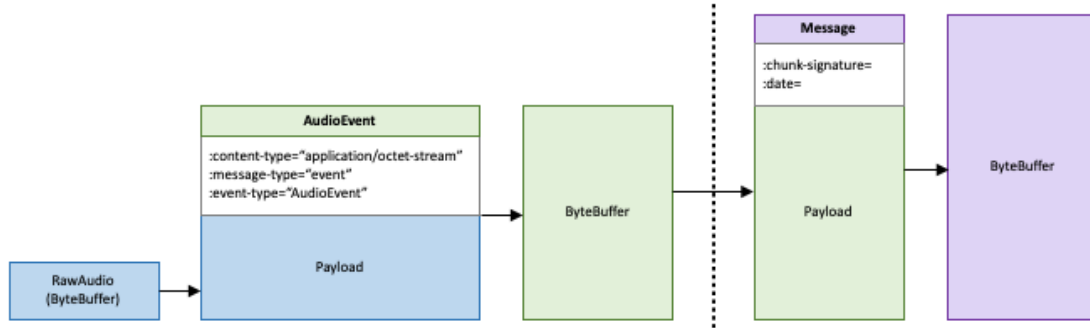
Data frames

Each request contains one or more data frames. The data frames use event stream encoding. The encoding supports bidirectional data transmission between a client and a server.

There are two steps to creating a data frame:

1. Combine the raw audio data with metadata to create the payload of the request.
2. Combine the payload with a signature to form the event message that is sent to Amazon Transcribe Medical.

The following diagram shows how this works.



Create the audio event

To create the message to send to Amazon Transcribe Medical, create the audio event. Combine the headers described in the following table with a chunk of audio bytes into an event-encoded message.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	24	application/octet-stream
11	:event-type	7	10	AudioEvent
13	:message-type	7	5	event

To create the payload for the event message, use a buffer in raw-byte format.

Create the message

Create a data frame using the audio event payload to send to Amazon Transcribe Medical. The data frame contains event-encoding headers that include the current date and a signature for the audio chunk and the audio event. To indicate to Amazon Transcribe Medical that the audio stream is complete, send an empty data frame that contains only the date and signature.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value
16	:chunk-signature	6	varies	generated signature
5	:date	8	8	timestamp

To create the signature for the data frame, first create a string to sign, and then calculate the signature for the event. Construct the string to sign as follows.

```
String stringToSign =
    "AWS4-HMAC-SHA256-PAYLOAD" +
    "\n" +
    DATE +
    "\n" +
    KEYPATH +
    "\n" +
    Hex(priorSignature) +
    "\n" +
```

```
HexHash(nonSignatureHeaders) +
"\n" +
HexHash(payload);
```

- **DATE:** The current date and time in Universal Time Coordinated (UTC) and using the [ISO 8601 format](#). Don't include milliseconds in the date. For example, 20220127T223754Z is 22:37:54 on 1/27/2022.
- **KEYPATH:** The signature scope in the format `date/region/service/aws4_request`. For example, `20220127/us-west-2/transcribe/aws4_request`.
- **priorSignature:** The signature for the previous frame. For the first data frame, use the signature of the header frame.
- **nonSignatureHeaders:** The `DATE` header encoded as a string.
- **payload:** The byte buffer containing the audio event data.
- **Hex:** A function that encodes its input into a hexadecimal representation.
- **HexHash:** A function that first creates a SHA-256 hash of its input and then uses the `Hex` function to encode the hash.

After you have constructed the string to sign, sign it using the key that you derived for Signature Version 4, as follows. For details, see [Examples of How to Derive a Signing Key for Signature Version 4](#) in the *AWS General Reference*.

```
String signature = HMACSHA256(derivedSigningKey, stringToSign);
```

- **HMACSHA256:** A function that creates a signature using the SHA-256 hash function.
- **derivedSigningKey:** The Signature Version 4 signing key.
- **stringToSign:** The string that you calculated for the data frame.

After you have calculated the signature for the data frame, construct a byte buffer containing the date, the signature, and the audio event payload. Send the byte array to Amazon Transcribe Medical for transcription.

End frame

To indicate that the audio stream is complete, send an end frame to Amazon Transcribe Medical. The *end frame* is a data frame with an empty payload. You construct the end frame the same way that you construct a data frame.

Streaming response

Responses from Amazon Transcribe Medical are also sent using event stream encoding. Use this information to decode a response from the [StartMedicalStreamTranscription](#) API.

Transcription response

A transcription response is event stream encoded. It contains the standard prelude and the following headers.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	16	application/json
11	:event-type	7	15	TranscriptEvent

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:message-type	7	5	event

For details, see [Event stream encoding \(p. 80\)](#).

When the response is decoded, it contains the following information:

```
:content-type: "application/vnd.amazon.eventstream"
:event-type: "TranscriptEvent"
:message-type: "event"
```

JSON transcription information

For an example of the JSON structure returned by Amazon Transcribe Medical, see [Using Amazon Transcribe Medical streaming with HTTP/2 \(p. 270\)](#).

Exception response

If there is an error in processing your transcription stream, Amazon Transcribe Medical sends an exception response. The response is event stream encoded. For details, see [Event stream encoding \(p. 80\)](#).

The response contains the standard prelude and the following headers.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	16	application/json
11	:event-type	7	19	BadRequestException
13	:message-type	7	9	exception

When the exception response is decoded, it contains the following information.

```
:content-type: "application/vnd.amazon.eventstream"
:event-type: "BadRequestException"
:message-type: "exception"
```

Exception message

Example request and response

The following is an end-to-end example of a streaming transcription request. In this example, binary data are represented as base64-encoded strings. In an actual response, the data are raw bytes.

Step 1: Start the session with Amazon Transcribe Medical

To start the session, send an HTTP/2 request to Amazon Transcribe Medical.

```
POST /medical-stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
```



```
authorization: Generated value
x-amz-target: com.amazonaws.transcribe.Transcribe.StartMedicalStreamTranscription
x-amz-content-sha256: STREAMING-MED-AWS4-HMAC-SHA256-EVENTS
x-amz-date: 20220208T235959Z
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
content-type: application/vnd.amazon.eventstream
transfer-encoding: chunked
```

Step 2: Send authentication information to Amazon Transcribe Medical

Amazon Transcribe Medical sends the following response.

```
HTTP/2.0 200
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-sample-rate: 16000
x-amzn-request-id: 8a08df7d-5998-48bf-a303-484355b4ab4e
x-amzn-transcribe-session-id: b4526fcf-5eee-4361-8192-d1cb9e9d6887
x-amzn-transcribe-media-encoding: flac
x-amzn-RequestId: 8a08df7d-5998-48bf-a303-484355b4ab4e
content-type: application/vnd.amazon.eventstream
```

Step 3: Create an audio event

Create an audio event containing the audio data to send. For details, see [Event stream encoding](#) (p. 269). The binary data in this request are base64-encoded. In an actual request, the data are raw bytes.

```
:content-type: "application/octet-stream"
:event-type: "AudioEvent"
:message-type: "event"

UklGRjzxPQBxQVZFZm10IBAAAAAABAAEAgD4AAAB9AAACABAAZGF0YVVTwPQAAAAAAAAAAAAAAAAAD//wIA/f8EAA==
```

Step 4: Create an audio event message

Create an audio message that contains the audio data to send to Amazon Transcribe Medical. For details, see [Event stream encoding](#) (p. 269). The audio event data in this example are base64-encoded. In an actual request, the data are raw bytes.

```
:date: 2019-01-29T01:56:17.291Z
:chunk-signature: signature

AAAA0gAAAIKVoRFcTTCjb250ZW50LXR5cGUHABhchBSaWNhdGlvbi9vY3RldC1zdHJlYW0LOmV2ZW50LXR5cGUHAApBdWRpb0V2ZW50DTptZXNzYWdlLXR5cGUHAAVldmVudAxDb256ZW50LVR5cGUHABphcHBSaWNhdGlvbi94LWFteilqc29uLTFEuMVJJRky88T0AV0FWRWZtdCAQAAAAAQABAIA+AAAAfQAAAQAQAGRhdGFu8D0AAAAA
AAAAAAAAAAAA//8CAP3/BAC7QLFf
```

Step 5: Use the response from Amazon Transcribe Medical

Amazon Transcribe Medical creates a stream of transcription events that it sends to your application. The events are sent in raw-byte format. In this example, the bytes are base64-encoded.

Amazon Transcribe Medical sends the following response.

```
AAAAUwAAAEPlRHpYBTpkYXRlCAAaAWiXUkMLEDPjaHVuay1zaWduYXR1cmUGACct6Zy+uymwEK2SrLp/zVBI  
5eGn83jdBwCaRUBJA+eaDafqjqI=
```

To see the transcription results, decode the raw bytes using event stream encoding.

```
:event-type: "TranscriptEvent"  
:content-type: "application/vnd.amazon.eventstream"  
:message-type: "event"  
  
{"Transcript":{"Results":[results]}}
```

For an example of the JSON structure returned by Amazon Transcribe Medical, see [Event stream encoding](#) (p. 269).

Step 6: End the transcription stream

Finally, send an empty audio event to Amazon Transcribe Medical to end the transcription stream. Create the audio event exactly like any other, except with an empty payload. Sign the event and include the signature in the `:chunk-signature` header, as follows.

```
:date: 2019-01-29T01:56:17.291Z  
:chunk-signature: signature
```

HTTP/2 streaming retry client

You can use the following code in your applications to handle retry logic for Amazon Transcribe Medical streaming transcription. The code provides tolerance for intermittent failures in the connection to Amazon Transcribe Medical. There are two parts of the client: an interface that you implement for your application, and the retry client itself.

Streaming retry client code

This code implements a streaming retry client. It manages the connection to Amazon Transcribe Medical and retries sending data when there are errors on the connection. For example, if there is a transient error on the network, this client resends the request that failed.

The retry client has two properties that control the behavior of the client. You can set:

- The maximum number of times that the client should attempt before failing. Reduce this value to make your application stop retrying sooner when there are network issues. The default is 10.
- The time in milliseconds that the client should wait between retries. Longer times raise the risk of losing data, shorter times raise the risk of your application being throttled. The default is 100 milliseconds.

The following is the client. You can copy this code to your application or use it as a starting point for your own client.

```
package com.amazonaws.wolverine.streaming;  
  
import com.amazonaws.transcribe.streaming.BadRequestException;
```

```
import com.amazonaws.transcribe.streaming.StartStreamTranscriptionRequest;
import org.reactivestreams.Publisher;
import org.slf4j.LoggerFactory;
import org.slf4j.Logger;
import software.amazon.awssdk.auth.credentials.EnvironmentVariableCredentialsProvider;
import software.amazon.awssdk.regions.Region;
import software.amazon.awssdk.auth.credentials.AwsCredentialsProvider;
import software.amazon.awssdk.auth.signer.EventStreamAws4Signer;
import software.amazon.awssdk.core.client.config.SdkAdvancedClientOption;
import software.amazon.awssdk.services.transcribestreaming.TranscribeStreamingAsyncClient;
import software.amazon.awssdk.services.transcribestreaming.model.AudioStream;
import
    software.amazon.awssdk.services.transcribestreaming.model.StartMedicalStreamTranscriptionRequest;
import
    software.amazon.awssdk.services.transcribestreaming.model.StartMedicalStreamTranscriptionResponseHandl

import java.net.URI;
import java.net.URISyntaxException;
import java.util.Arrays;
import java.util.List;
import java.util.UUID;
import java.util.concurrent.CompletableFuture;

public class TranscribeStreamingRetryClient {

    private static final int DEFAULT_MAX_RETRIES = 10;
    private static final int DEFAULT_MAX_SLEEP_TIME_MILLS = 100;
    private static final Logger log =
        LoggerFactory.getLogger(TranscribeStreamingRetryClient.class);
    private final TranscribeStreamingAsyncClient client;
    List<Class<?>> nonRetriableExceptions = Arrays.asList(BadRequestException.class);
    private int maxRetries = DEFAULT_MAX_RETRIES;
    private int sleepTime = DEFAULT_MAX_SLEEP_TIME_MILLS;

    /**
     * Create a TranscribeStreamingRetryClient with given credential and configuration
     */
    public TranscribeStreamingRetryClient(AwsCredentialsProvider creds,
                                           String endpoint, Region region) throws
        URISyntaxException {
        this(TranscribeStreamingAsyncClient.builder()
            .overrideConfiguration(
                c -> c.putAdvancedOption(
                    SdkAdvancedClientOption.SIGNER,
                    EventStreamAws4Signer.create()))
            .credentialsProvider(creds)
            .endpointOverride(new URI(endpoint))
            .region(region)
            .build());
    }

    /**
     * Initiate TranscribeStreamingRetryClient with TranscribeStreamingAsyncClient
     */
    public TranscribeStreamingRetryClient(TranscribeStreamingAsyncClient client) {
        this.client = client;
    }

    /**
     * Get Max retries
     */
    public int getMaxRetries() {
        return maxRetries;
    }
}
```

```
/**
 * Set Max retries
 */
public void setMaxRetries(int maxRetries) {
    this.maxRetries = maxRetries;
}

/**
 * Get sleep time
 */
public int getSleepTime() {
    return sleepTime;
}

/**
 * Set sleep time between retries
 */
public void setSleepTime(int sleepTime) {
    this.sleepTime = sleepTime;
}

/**
 * Initiate a Stream Transcription with retry.
 */

    public CompletableFuture<Void> startStreamTranscription(final
StartStreamTranscriptionRequest request,
                                                    final Publisher<AudioStream>
publisher,
                                                    final
StreamTranscriptionBehavior responseHandler) {

        CompletableFuture<Void> finalFuture = new CompletableFuture<>();

        recursiveStartStream(rebuildRequestWithSession(request), publisher,
responseHandler, finalFuture, 0);

        return finalFuture;
    }

/**
 * Recursively call startStreamTranscription() until the request is completed or we run
out of retries.
 */
private void recursiveStartStream(final StartMedicalStreamTranscriptionRequest request,
final Publisher<AudioStream> publisher,
final StreamTranscriptionBehavior responseHandler,
final CompletableFuture<Void> finalFuture,
final int retryAttempt) {
    StartMedicalStreamTranscriptionResponseHandler responseHandler1 =
getResponseHandler(responseHandler);
    CompletableFuture<Void> result = client.startMedicalStreamTranscription(request,
publisher, responseHandler1);
    result.whenComplete((r, e) -> {
        if (e != null) {
            log.debug("Error occurred:", e);

            if (retryAttempt <= maxRetries && isExceptionRetriable(e)) {
                log.debug("Retriable error occurred and will be retried.");
                log.debug("Sleeping for sometime before retrying...");
                try {
                    Thread.sleep(sleepTime);
                } catch (InterruptedException e1) {
                    log.debug("Unable to sleep. Failed with exception: ", e);
                    e1.printStackTrace();
                }
            }
        }
    });
}
```

```
        log.debug("Making retry attempt: " + (retryAttempt + 1));
        recursiveStartStream(request, publisher, responseHandler, finalFuture,
retryAttempt + 1);
    } else {
        log.error("Encountered unretriable exception or ran out of retries. ");
        responseHandler.onError(e);
        finalFuture.completeExceptionally(e);
    }
} else {
    responseHandler.onComplete();
    finalFuture.complete(null);
}
});
}

private StartMedicalStreamTranscriptionRequest
rebuildRequestWithSession(StartStreamTranscriptionRequest request) {
    return StartMedicalStreamTranscriptionRequest.builder()
        .languageCode(request.getLanguageCode())
        .mediaEncoding(request.getMediaEncoding())
        .mediaSampleRateHertz(request.getMediaSampleRateHertz())
        .sessionId(UUID.randomUUID().toString())
        .build();
}

/**
 * StartStreamTranscriptionResponseHandler implements subscriber of transcript stream
 * Output is printed to standard output
 */
private StartMedicalStreamTranscriptionResponseHandler getResponseHandler(
    StreamTranscriptionBehavior transcriptionBehavior) {
    final StartMedicalStreamTranscriptionResponseHandler build =
StartMedicalStreamTranscriptionResponseHandler.builder()
        .onResponse(transcriptionBehavior::onResponse)
        .onError(e -> {
            //Do nothing here. Don't close any streams that shouldn't be cleaned up
yet.
        })
        .onComplete(() -> {
            //Do nothing here. Don't close any streams that shouldn't be cleaned up
yet.
        })
        .subscriber(transcriptionBehavior::onStream)
        .build();
    return build;
}

/**
 * Check if the exception can be retried.
 */
private boolean isExceptionRetriable(Throwable e) {
    e.printStackTrace();

    return nonRetriableExceptions.contains(e.getClass());
}

public void close() {
    this.client.close();
}

public static void main(String[] args) throws URISyntaxException {
    TranscribeStreamingRetryClient transcribeStreamingRetryClient = new
TranscribeStreamingRetryClient(EnvironmentVariableCredentialsProvider.create(),
        "https://transcribestreaming.us-west-2.amazonaws.com", Region.US_WEST_2);
//    transcribeStreamingRetryClient.startStreamTranscription(...)
```

```
}  
}
```

Streaming retry client interface code

This interface is similar to the response handler used in the getting started example. It implements the same event handlers. Implement this interface to use the streaming-med retry client.

```
package com.amazonaws.wolverine.streaming;  
  
import  
    software.amazon.awssdk.services.transcribestreaming.model.MedicalTranscriptResultStream;  
import  
    software.amazon.awssdk.services.transcribestreaming.model.StartMedicalStreamTranscriptionResponse;  
import  
    software.amazon.awssdk.services.transcribestreaming.model.StartStreamTranscriptionResponse;  
import software.amazon.awssdk.services.transcribestreaming.model.TranscriptResultStream;  
  
/**  
 * Defines how a stream response should be handled.  
 * You should build a class implementing this interface to define the behavior.  
 */  
public interface StreamTranscriptionBehavior {  
    /**  
     * Defines how to respond when encountering an error on the stream transcription.  
     */  
    void onError(Throwable e);  
  
    /**  
     * Defines how to respond to the Transcript result stream.  
     */  
    void onStream(MedicalTranscriptResultStream e);  
  
    /**  
     * Defines what to do on initiating a stream connection with the service.  
     */  
    void onResponse(StartMedicalStreamTranscriptionResponse r);  
  
    /**  
     * Defines what to do on stream completion  
     */  
    void onComplete();  
}
```

The following is an example implementation of the `StreamTranscriptionBehavior` interface. You can use this implementation or use it as a starting point for your own implementation.

```
package com.amazonaws.wolverine.streaming;  
  
import software.amazon.awssdk.services.transcribestreaming.model.MedicalResult;  
import software.amazon.awssdk.services.transcribestreaming.model.MedicalTranscriptEvent;  
import  
    software.amazon.awssdk.services.transcribestreaming.model.MedicalTranscriptResultStream;  
import  
    software.amazon.awssdk.services.transcribestreaming.model.StartMedicalStreamTranscriptionResponse;  
  
import java.util.List;  
  
/**  
 * Implementation of StreamTranscriptionBehavior to define how a stream response should be  
 * handled.  
 */
```

```
*
* COPYRIGHT:
*
* Copyright 2018 Amazon.com, Inc. or its affiliates. All Rights Reserved.
*
* Licensed under the Apache License, Version 2.0 (the "License").
* You may not use this file except in compliance with the License.
* A copy of the License is located at
*
*     http://www.apache.org/licenses/LICENSE-2.0
*
* or in the "license" file accompanying this file. This file is distributed
* on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either
* express or implied. See the License for the specific language governing
* permissions and limitations under the License.
*/
public class StreamTranscriptionBehaviorImpl implements StreamTranscriptionBehavior {

    @Override
    public void onError(Throwable e) {
        System.out.println("=== Failure Encountered ===");
        e.printStackTrace();
    }

    @Override
    public void onStream(MedicalTranscriptResultStream e) {
        // EventResultStream has other fields related to the timestamp of the transcripts
        in it.
        // Please refer to the javadoc of TranscriptResultStream for more details
        List<MedicalResult> results = ((MedicalTranscriptEvent) e).transcript().results();
        if (results.size() > 0) {
            if (results.get(0).alternatives().size() > 0)
                if (!results.get(0).alternatives().get(0).transcript().isEmpty()) {
                    System.out.println(results.get(0).alternatives().get(0).transcript());
                }
        }
    }

    @Override
    public void onResponse(StartMedicalStreamTranscriptionResponse r) {
        System.out.println(String.format("=== Received Initial response. Request Id: %s
        ===", r.requestId()));
    }

    @Override
    public void onComplete() {
        System.out.println("=== All records stream successfully ===");
    }
}
```

Next step

[Using the HTTP/2 retry client \(p. 281\)](#)

Using the HTTP/2 retry client

The following is a sample application that uses the retry client to transcribe audio from either a file or your microphone. You can use this application to test the client, or you can use it as a starting point for your own applications.

To run the sample, do the following:

- Copy the retry client to your workspace. See [Streaming retry client code \(p. 276\)](#).
- Copy the retry client interface to your workspace. Implement the interface, or you can use the sample implementation. See [Streaming retry client interface code \(p. 280\)](#).
- Copy the sample application to your workspace. Build and run the application.

```
public class StreamingRetryApp {
    private static final String endpoint = "endpoint";
    private static final Region region = Region.US_EAST_1;
    private static final int sample_rate = 28800;
    private static final String encoding = " ";
    private static final String language = LanguageCode.EN_US.toString();

    public static void main(String args[]) throws URISyntaxException, ExecutionException,
        InterruptedException, LineUnavailableException, FileNotFoundException {
        /**
         * Create Amazon Transcribe streaming retry client.
         */

        TranscribeStreamingRetryClient client = new
        TranscribeStreamingRetryClient(EnvironmentVariableCredentialsProvider.create(), endpoint,
        region);

        StartStreamTranscriptionRequest request = StartStreamTranscriptionRequest.builder()
            .languageCode(language)
            .mediaEncoding(encoding)
            .mediaSampleRateHertz(sample_rate)
            .build();

        /**
         * Start real-time speech recognition. The Amazon Transcribe streaming java client
         uses the Reactive-streams
         * interface. For reference on Reactive-streams:
         * https://github.com/reactive-streams/reactive-streams-jvm
         */
        CompletableFuture<Void> result = client.startStreamTranscription(
            /**
             * Request parameters. Refer to API documentation for details.
             */
            request,
            /**
             * Provide an input audio stream.
             * For input from a microphone, use getStreamFromMic().
             * For input from a file, use getStreamFromFile().
             */
            new AudioStreamPublisher(
                new FileInputStream(new File("FileName"))),
            /**
             * Object that defines the behavior on how to handle the stream
             */
            new StreamTranscriptionBehaviorImpl());

        /**
         * Synchronous wait for stream to close, and close client connection
         */
        result.get();
        client.close();
    }

    private static class AudioStreamPublisher implements Publisher<AudioStream> {
        private final InputStream inputStream;
        private static Subscription currentSubscription;

        private AudioStreamPublisher(InputStream inputStream) {
            this.inputStream = inputStream;
        }
    }
}
```



```
    }  
  
    @Override  
    public void subscribe(Subscriber<? super AudioStream> s) {  
        if (this.currentSubscription == null) {  
            this.currentSubscription = new  
TranscribeStreamingDemoApp.SubscriptionImpl(s, inputStream);  
        } else {  
            this.currentSubscription.cancel();  
            this.currentSubscription = new  
TranscribeStreamingDemoApp.SubscriptionImpl(s, inputStream);  
        }  
        s.onSubscribe(currentSubscription);  
    }  
}  
}
```

Establish a bi-directional connection using the WebSocket protocol

You can use the [WebSocket protocol](#) to open a bi-directional connection that is designed to be secure with Amazon Transcribe Medical. You encode the audio stream with event stream encoding, and Amazon Transcribe Medical returns a stream of text in JSON blobs that are also encoded using event stream encoding. For more information, see [Event stream encoding \(p. 269\)](#). You can use the information in this section to create applications using the WebSocket library of your choice.

Topics

- [Adding a policy for WebSocket requests to your IAM role \(p. 283\)](#)
- [Creating a pre-signed URI \(p. 284\)](#)
- [Handling the WebSocket upgrade response \(p. 288\)](#)
- [Making a WebSocket streaming request \(p. 288\)](#)
- [Handling a WebSocket streaming response \(p. 289\)](#)
- [Handling WebSocket streaming errors \(p. 289\)](#)

Adding a policy for WebSocket requests to your IAM role

To use the WebSocket protocol to call Amazon Transcribe Medical, you need to attach the following policy to the IAM role that makes the request.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Sid": "transcribemedicalstreaming",  
      "Effect": "Allow",  
      "Action": "transcribe:StartMedicalStreamTranscription",  
      "Resource": "*"   
    }  
  ]  
}
```

Creating a pre-signed URI

You need to construct a URI for your WebSocket request that contains the information needed to set up communication between your application and Amazon Transcribe Medical. To open a bi-directional connection, you must use port 8443. WebSocket streaming-med uses the Amazon Signature Version 4 process for signing requests. Signing the request helps to verify the identity of the requester and to protect your audio data in transit. It also protects against potential replay attacks. For more information about Signature Version 4, see [Signing AWS API Requests](#) in the *Amazon Web Services General Reference*.

The URI has the following format. Line breaks have been added for readability.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/medical-stream-transcription-  
websocket  
?language-code=languageCode  
&X-Amz-Algorithm=AWS4-HMAC-SHA256  
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request  
&X-Amz-Date=20220208T235959Z  
&X-Amz-Expires=250  
&X-Amz-Security-Token=security-token  
&X-Amz-Signature=Signature Version 4 signature  
&X-Amz-SignedHeaders=host  
&media-encoding=flac  
&sample-rate=16000  
&session-id=sessionId  
&specialty=specialty  
&type=type
```

For TYPE, it is best to select CONVERSATION if your use case involves multiple speakers engaged in a discussion. An example would be a conversation between a clinician and a patient. Select DICTATION if your use case involves a single speaker where a person is dictating speech. An example would be if a physician is dictating medical notes for data entry purposes after a patient encounter.

Use the following values for the URI parameters:

- **language-code** – The language code for the input audio. The valid value is en-US.
- **media-encoding** – The encoding used for the input audio. The valid value is pcm.
- **sample-rate** – The sample rate of the input audio in hertz; sample rates of 16,000 Hz or higher are accepted.
- **sessionId** – Optional. An identifier for the transcription session. If you don't provide a session ID, Amazon Transcribe Medical generates one for you and returns it in the response.
- **specialty** – The specialty of the medical domain. Valid values are PRIMARYCARE, CARDIOLOGY, NEUROLOGY, ONCOLOGY, RADIOLOGY, and UROLOGY.
- **type** – The type of audio. Must be DICTATION or CONVERSATION.

The remaining parameters are Signature Version 4 parameters:

- **X-Amz-Algorithm** – The algorithm you're using in the signing process. The only valid value is AWS4-HMAC-SHA256.
- **X-Amz-Credential** – A string formed by concatenating your access key ID, the date, your AWS Region, the AWS service you're calling, and a standard termination term. Each of these components is separated by '%2F'. This is the general form:

```
your-access-key-id%2Fdate%2FAWS-Region%2FAWS-service%2Faws4_request
```

For more information, refer to [Handling Dates in Signature Version 4](#).

- **X-Amz-Date** – The date and time the signature is created. The format is YYYYMMDDTHHMMSSZ, where YYYY=year, MM=month, DD=day, HH=hour, MM=minute, SS=seconds, and 'T' and 'Z' are fixed characters. For more information, refer to [Handling Dates in Signature Version 4](#).
- **X-Amz-Expires** – The length of time in seconds until the credentials expire. The maximum value is 300 seconds (5 minutes).
- **X-Amz-Security-Token** – Optional. A Signature Version 4 token for temporary credentials. If you specify this parameter, include it in the canonical request. For more information, see [Requesting Temporary Security Credentials](#) in the *AWS Identity and Access Management User Guide*.
- **X-Amz-Signature** – The Signature Version 4 signature that you generated for the request.
- **X-Amz-SignedHeaders** – The headers that are signed when creating the signature for the request. The only valid value is host.

To construct the URI for the request and create the Signature Version 4 signature, use the following steps. The examples are in pseudocode.

Task 1: Create a canonical request

Create a string that includes information from your request in a standardized format. This ensures that when AWS receives the request, it can calculate the same signature that you calculate in Task 3. For more information, see [Create a Canonical Request for Signature Version 4](#) in the *AWS General Reference*.

1. Define variables for the request in your application.

```
# HTTP verb
method = "GET"
# Service name
service = "transcribe"
# Region
region = "us-west-2"
# Amazon Transcribe streaming-med endpoint
endpoint = "https://transcribestreaming.us-west-2.amazonaws.com:8443"
# Host
host = "transcribestreaming.us-west-2.amazonaws.com:8443"
# Date and time of request
amz-date = "YYYYMMDD'T'HHMMSS'Z'"
# Date without time for credential scope
datestamp = "YYYYMMDD"
```

2. Create a canonical URI. The canonical URI is the part of the URI between the domain and the query string.

```
canonical_uri = "/medical-stream-transcription-websocket"
```

3. Create the canonical headers and signed headers. Note the trailing "\n" in the canonical headers.

- Append the lowercase header name followed by a colon.
- Append a comma-separated list of values for that header. Do not sort the values in headers that have multiple values.
- Append a new line (\n).

```
canonical_headers = "host:" + host + "\n"
signed_headers = "host"
```

4. Match the algorithm to the hashing algorithm. You must use SHA-256.

```
algorithm = "AWS4-HMAC-SHA256"
```

5. Create the credential scope, which scopes the derived key to the date, AWS Region, and service to which the request is made.

```
credential_scope = datestamp + "/" + region + "/" + service + "/" + "aws4_request"
```

6. Create the canonical query string. Query string values must be URI-encoded and sorted by name.
 - Sort the parameter names by character code point in ascending order. Parameters with duplicate names should be sorted by value. For example, a parameter name that begins with the uppercase letter F precedes a parameter name that begins with a lowercase letter b.
 - Do not URI-encode any of the unreserved characters that [RFC 3986](#) defines: A-Z, a-z, 0-9, hyphen (-), underscore (_), period (.), and tilde (~).
 - Percent-encode all other characters with %XY, where X and Y are hexadecimal characters (0-9 and uppercase A-F). For example, the space character must be encoded as %20 (not using '+', as some encoding schemes do) and extended UTF-8 characters must be in the form %XY%ZA%BC.
 - Double-encode any equals (=) characters in parameter values.

```
canonical_querystring = "X-Amz-Algorithm=" + algorithm
canonical_querystring += "&X-Amz-Credential=" + URI-encode(access key + "/" +
  credential_scope)
canonical_querystring += "&X-Amz-Date=" + amz_date
canonical_querystring += "&X-Amz-Expires=250"
canonical_querystring += "&X-Amz-Security-Token=" + token
canonical_querystring += "&X-Amz-SignedHeaders=" + signed_headers
canonical_querystring += "&language-code=en-US&media-encoding=flac&sample-rate=16000"
canonical_querystring += "&specialty=PRIMARYCARE"
canonical_querystring += "&type=Dictation"
```

7. Create a hash of the payload. For a GET request, the payload is an empty string.

```
payload_hash = HashSHA256("").Encode("utf-8").HexDigest()
```

8. Combine all of the elements to create the canonical request.

```
canonical_request = method + '\n'
  + canonical_uri + '\n'
  + canonical_querystring + '\n'
  + canonical_headers + '\n'
  + signed_headers + '\n'
  + payload_hash
```

Task 2: Create the string to sign

The string to sign contains meta information about your request. You use the string to sign in the next step when you calculate the request signature. For more information, see [Create a String to Sign for Signature Version 4](#) in the *AWS General Reference*.

- Create the string.

```
string_to_sign=algorithm + "\n"
  + amz_date + "\n"
  + credential_scope + "\n"
```

```
+ HashSHA256(canonical_request.Encode("utf-8").HexDigest())
```

Task 3: Calculate the signature

You derive a signing key from your AWS secret access key. For a greater degree of protection, the derived key is specific to the date, service, and AWS Region. You use the derived key to sign the request. For more information, see [Calculate the Signature for AWS Signature Version 4](#) in the *AWS General Reference*.

The code assumes that you have implemented the `GetSignatureKey` function to derive a signing key. For more information and example functions, see [Examples of How to Derive a Signing Key for Signature Version 4](#) in the *AWS General Reference*.

The function `HMAC(key, data)` represents an HMAC-SHA256 function that returns the results in binary format.

- Create the signing key and sign the string to sign.

```
#Create the signing key
signing_key = GetSignatureKey(secret_key, timestamp, region, service)

# Sign the string_to_sign using the signing key
signature = HMAC.new(signing_key, (string_to_sign).Encode("utf-8"), Sha256()).HexDigest
```

Task 4: Add signing information to the request and create the request URI

After you calculate the signature, add it to the query string. For more information, see [Add the Signature to the HTTP Request](#) in the *AWS General Reference*.

1. Add the authentication information to the query string.

```
canonical_querystring += "&X-Amz-Signature=" + signature
```

2. Create the URI for the request.

```
request_url = endpoint + canonical_uri + "?" + canonical_querystring
```

You use the request URI with your WebSocket library to make the request to the Amazon Transcribe Medical service.

Including WebSocket request headers

The request to Amazon Transcribe Medical must include the following headers. Typically these headers are managed by your WebSocket client library.

```
Host: transcribestreaming.us-west-2.amazonaws.com:8443
Connection: Upgrade
Upgrade: websocket
Origin: request source
Sec-WebSocket-Version: 13
Sec-WebSocket-Key: random key
```

Use the following values for the headers:

- **Connection** – Always Upgrade.

- **Upgrade** – Always websocket.
- **Origin** – The URI of the WebSocket client.
- **Sec-WebSocket-Version** – The version of the WebSocket protocol to use.
- **Sec-WebSocket-Key** – A base-64 encoded randomly generated string that identifies the request.

Handling the WebSocket upgrade response

When Amazon Transcribe Medical receives your WebSocket request, it responds with a WebSocket upgrade response. Typically, your WebSocket library manages this response and sets up a socket for communications with Amazon Transcribe Medical.

The following is the response from Amazon Transcribe Medical. Line breaks have been added to the `websocket-location` header for readability.

```
HTTP/1.1 101 WebSocket Protocol Handshake
Connection: upgrade
Upgrade: websocket
websocket-origin: https://transcribestreaming.us-west-2.amazonaws.com:8443
websocket-location: transcribestreaming.us-west-2.amazonaws.com:8443/medical-stream-
transcription-websocket?
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-SignedHeaders=host
&language-code=en-US
&media-encoding=flac
&sample-rate=16000
&type=dictation or conversation
&specialty=medical specialty (must be Primary Care)
&X-Amz-Signature=signature
x-amzn-RequestId: RequestId
x-amzn-SessionId: SessionId
Strict-Transport-Security: max-age=31536000
sec-websocket-accept: token
```

The response has the following values:

- **Connection** – Always Upgrade.
- **Upgrade** – Always websocket.
- **websocket-origin** – The URI of the WebSocket server that responded to the request.
- **websocket-location** – The contents of the request URI that was sent to the server.
- **x-amzn-RequestId** – An identifier for the request.
- **x-amzn-SessionId** – An identifier for a transcription session.
- **Strict-Transport-Security** – A header that informs browsers to access the endpoint using only HTTPS.
- **sec-websocket-accept** – The hash of the `Sec-WebSocket-Key` header sent in the request.

Making a WebSocket streaming request

After the WebSocket connection is established, the client can start sending a sequence of audio frames. Each frame contains one data frame that is encoded in event stream encoding. For more information, see [Event stream encoding](#) (p. 80).

Each data frame contains three headers combined with a chunk of raw audio bytes. The following table lists and describes the headers.

Header name Byte Length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	24	application/ octet-stream
11	:event-type	7	10	AudioEvent
13	:message-type	7	5	event

To end the audio data stream, send an empty audio chunk in an event stream encoded message.

Handling a WebSocket streaming response

The response contains event stream encoded raw bytes in the payload. It contains the standard prelude and the following headers.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	24	application/octet- stream
11	:event-type	7	15	TranscriptEvent
13	:message-type	7	5	event

When you decode the binary response, you end up with a JSON structure with the results of the transcription. For an example of the JSON response, see [Streaming transcription \(p. 264\)](#).

Handling WebSocket streaming errors

If an exception occurs while processing your request, Amazon Transcribe Medical responds with a terminal WebSocket frame containing an event stream encoded response. The response has the headers described in the following table, and the body of the response contains a descriptive error message. After sending the exception response, Amazon Transcribe Medical sends a close frame.

Header name byte length	Header name (string)	Header value type	Value string byte length	Value string (UTF-8)
13	:content-type	7	24	application/octet- stream
15	:exception- type	7	varies	varies, see below
13	:message-type	7	9	exception

The `exception-type` header contains one of the following values:

- **BadRequestException** – There was a client error when the stream was created, or an error occurred while streaming-med data. Make sure that your client is ready to accept data and try your request again.
- **InternalFailureException** – Amazon Transcribe Medical had a problem during the handshake with the client. Try your request again.
- **LimitExceededException** – The client exceeded the concurrent stream limit. For more information, see [Amazon Transcribe Service Quotas](#) in the *Amazon Web Services General Reference*. Reduce the number of streams that you are transcribing.
- **UnrecognizedClientException** – The WebSocket upgrade request was signed with an incorrect access key or secret key. Make sure that you are correctly creating the access key and try your request again.

In addition, Transcribe Medical can return any of the common service errors. For a list, see [Common Errors](#).

Transcribing a medical conversation

You can use Amazon Transcribe Medical to transcribe a medical conversation between a clinician and a patient using either a batch transcription job or a real-time stream. Batch transcription jobs enable you to transcribe audio files. To ensure that Amazon Transcribe Medical produces transcription results with the highest possible accuracy, you must specify the medical specialty of the clinician in your transcription job or stream.

You can transcribe a clinician-patient visit in the following medical specialties:

- Cardiology – available in streaming transcription only
- Neurology – available in streaming transcription only
- Oncology – available in streaming transcription only
- Primary Care – includes the following types of medical practice:
 - Family medicine
 - Internal medicine
 - Obstetrics and Gynecology (OB-GYN)
 - Pediatrics
- Urology – available in streaming transcription only

You can improve transcription accuracy by using medical custom vocabularies. For information on how medical custom vocabularies work, see [Improving transcription accuracy with medical custom vocabularies \(p. 313\)](#).

By default, Amazon Transcribe Medical returns the transcription with the highest confidence level. If you'd like to configure it to return alternative transcriptions, see [Generating alternative transcriptions \(p. 327\)](#).

For information about how numbers and medical measurements appear in the transcription output, see [Transcribing numbers \(p. 259\)](#) and [Transcribing medical terms and measurements \(p. 260\)](#).

Topics

- [Transcribing an audio file of a medical conversation \(p. 291\)](#)
- [Transcribing a medical conversation in a real-time stream \(p. 294\)](#)
- [Identifying speakers and labeling their speech \(p. 296\)](#)
- [Transcribing multi-channel audio \(p. 302\)](#)

Transcribing an audio file of a medical conversation

Use a batch transcription job to transcribe audio files of medical conversations. You can use this to transcribe a clinician-patient dialogue. You can start a batch transcription job in either the [StartMedicalTranscriptionJob](#) API or the AWS Management Console.

When you start a medical transcription job with the [StartMedicalTranscriptionJob](#) API, you specify `PRIMARYCARE` as the value of the `Specialty` parameter.

AWS Management Console

To transcribe a clinician-patient dialogue (AWS Management Console)

To use the AWS Management Console to transcribe a clinician-patient dialogue, create a transcription job and choose **Conversation** for **Audio input type**.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, under **Job settings**, specify the following.
 - a. **Name** – the name of the transcription job.
 - b. **Audio input type** – **Conversation**
5. For the remaining fields, specify the Amazon S3 location of your audio file and where you want to store the output of your transcription job.
6. Choose **Next**.
7. Choose **Create**.

API

To transcribe a medical conversation using a batch transcription job (API)

- For the [StartMedicalTranscriptionJob](#) API, specify the following.
 - a. For `MedicalTranscriptionJobName`, specify a name unique in your AWS account.
 - b. For `LanguageCode`, specify the language code that corresponds to the language spoken in your audio file and the language of your vocabulary filter.
 - c. For the `MediaFileUri` parameter of the `Media` object, specify the name of the audio file that you want to transcribe.
 - d. For `Specialty`, specify the medical specialty of the clinician speaking in the audio file as `PRIMARYCARE`.
 - e. For `Type`, specify `CONVERSATION`.
 - f. For `OutputBucketName`, specify the Amazon S3 bucket to store the transcription results.

The following is an example request that uses the AWS SDK for Python (Boto3) to transcribe a medical conversation of a clinician in the `PRIMARYCARE` specialty and a patient.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-med-transcription-job"
```

```
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac"
transcribe.start_medical_transcription_job(
    MedicalTranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'output-files/',
    LanguageCode = 'en-US',
    Specialty = 'PRIMARYCARE',
    Type = 'CONVERSATION'
)

while True:
    status = transcribe.get_medical_transcription_job(MedicalTranscriptionJobName =
    job_name)
    if status['MedicalTranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED',
    'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

The following example code shows the transcription results of a clinician-patient conversation.

```
{
  "jobName": "conversation-medical-transcription-job",
  "accountId": "111122223333",
  "results": {
    "transcripts": [
      {
        "transcript": "... come for a follow up visit today..."
      }
    ],
    "items": [
      {
        ...
        "start_time": "4.85",
        "end_time": "5.12",
        "alternatives": [
          {
            "confidence": "1.0",
            "content": "come"
          }
        ],
        "type": "pronunciation"
      },
      {
        "start_time": "5.12",
        "end_time": "5.29",
        "alternatives": [
          {
            "confidence": "1.0",
            "content": "for"
          }
        ],
        "type": "pronunciation"
      },
      {
        "start_time": "5.29",
        "end_time": "5.33",
```

```
    "alternatives": [
      {
        "confidence": "0.9955",
        "content": "a"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "5.33",
    "end_time": "5.66",
    "alternatives": [
      {
        "confidence": "0.9754",
        "content": "follow"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "5.66",
    "end_time": "5.75",
    "alternatives": [
      {
        "confidence": "0.9754",
        "content": "up"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "5.75",
    "end_time": "6.02",
    "alternatives": [
      {
        "confidence": "1.0",
        "content": "visit"
      }
    ]
  }
  ...
},
"status": "COMPLETED"
}
```

AWS CLI

To transcribe a medical conversation using a batch transcription job (AWS CLI)

- Run the following code.

```
aws transcribe start-medical-transcription-job \  
--region us-west-2 \  
--cli-input-json file://example-start-command.json
```

The following code shows the contents of `example-start-command.json`.

```
{  
  "MedicalTranscriptionJobName": "my-first-med-transcription-job",  
  "Media": {
```

```
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "LanguageCode": "en-US",  
  "Specialty": "PRIMARYCARE",  
  "Type": "CONVERSATION"  
}
```

Transcribing a medical conversation in a real-time stream

You can transcribe an audio stream of a medical conversation using either the HTTP/2 or [WebSocket](#) protocols. For information on how to start a stream using the WebSocket protocol, see [Establish a bi-directional connection using the WebSocket protocol \(p. 283\)](#). To start an HTTP/2 stream, use the [StartMedicalStreamTranscription](#) API.

You can transcribe streaming audio in the following medical specialties:

- Cardiology
- Neurology
- Oncology
- Primary Care
- Urology

Each medical specialty includes many types of procedures and appointments. Clinicians therefore dictate many different types of notes. Use the following examples as guidance to help you specify the value of the `specialty` URI parameter of the WebSocket request, or the `Specialty` parameter of the [StartMedicalStreamTranscription](#) API:

- For electrophysiology or echocardiography consultations, choose `CARDIOLOGY`.
- For medical oncology, surgical oncology, or radiation oncology consultations, choose `ONCOLOGY`.
- For a physician providing a consultation to a patient who had a stroke, either a transient ischemic attack or a cerebrovascular attack, choose `NEUROLOGY`.
- For a consultation around urinary incontinence, choose `UROLOGY`.
- For yearly checkup or urgent care visits, choose `PRIMARYCARE`.
- For inpatient hospitalist visits, choose `PRIMARYCARE`.
- For consultations regarding fertility, tubal ligation, IUD insertion, or abortion, choose `PRIMARYCARE`.

AWS Management Console

To transcribe a streaming medical conversation (AWS Management Console)

To use the AWS Management Console to transcribe a clinician-patient dialogue in real-time stream, choose the option to transcribe a medical conversation, start the stream, and begin speaking into the microphone.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Real-time transcription**.
3. Choose **Conversation**.

4. For **Medical specialty**, choose the clinician's specialty.
5. Choose **Start streaming**.
6. Speak into the microphone.

Transcribing a medical conversation in an HTTP/2 stream

The following is the syntax for the parameters of an HTTP/2 request.

To transcribe an HTTP/2 stream of a medical conversation, use the [StartMedicalStreamTranscription](#) API and specify the following:

- `LanguageCode` – The language code. The valid value is `en-US`
- `MediaEncoding` – The encoding used for the input audio. Valid values are `pcm`, `ogg-opus`, and `flac`.
- `Specialty` – The specialty of the medical professional.
- `Type` – `CONVERSATION`

To improve transcription accuracy of specific terms in a real-time stream, use a custom vocabulary. To enable a custom vocabulary, set the value of `VocabularyName` parameter to the name of the custom vocabulary that you want to use. For more information, see [Improving transcription accuracy with medical custom vocabularies](#) (p. 313).

To label the speech from different speakers, set the `ShowSpeakerLabel` parameter to `true`. For more information, see [Identifying speakers and labeling their speech](#) (p. 296).

For more information on setting up an HTTP/2 stream to transcribe a medical conversation, see [Streaming request](#) (p. 270).

Transcribing a medical conversation in a WebSocket stream

You can use a WebSocket request to transcribe a medical conversation. When you make a WebSocket request, you create a pre-signed URI. This URI contains the information needed to set up the audio stream between your application and Amazon Transcribe Medical. For more information on creating WebSocket requests, see [Establish a bi-directional connection using the WebSocket protocol](#) (p. 283).

Use the following template to create your pre-signed URI.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/medical-stream-transcription-  
websocket  
?language-code=LanguageCode  
&X-Amz-Algorithm=AWS4-HMAC-SHA256  
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request  
&X-Amz-Date=20220208T235959Z  
&X-Amz-Expires=250  
&X-Amz-Security-Token=security-token  
&X-Amz-Signature=Signature Version 4 signature  
&X-Amz-SignedHeaders=host  
&media-encoding=flac  
&sample-rate=16000  
&session-id=sessionId  
&specialty=medicalSpecialty  
&type=CONVERSATION  
&vocabulary-name=vocabularyName  
&show-speaker-label=boolean
```

To improve transcription accuracy of specific terms in a real-time stream, use a custom vocabulary. To enable a custom vocabulary, set the value of `vocabulary-name` to the name of the custom vocabulary

that you want to use. For more information, see [Improving transcription accuracy with medical custom vocabularies \(p. 313\)](#).

To label the speech from different speakers, set the `show-speaker-label` parameter in to `true`. For more information, see [Identifying speakers and labeling their speech \(p. 296\)](#).

For more information on creating pre-signed URIs, see [Creating a pre-signed URI \(p. 284\)](#).

Identifying speakers and labeling their speech

To identify and label the speech from different speakers in Amazon Transcribe Medical, use *speaker diarization*. This enables you to see what the patient said and what the clinician said in the transcription output.

When you enable speaker diarization, Amazon Transcribe Medical labels each speaker *utterance* with a unique identifier for each speaker. An *utterance* is a unit of speech that is typically separated from other utterances by silence. In batch transcription, an utterance from the clinician could receive a label of `spk_0` and an utterance the patient could receive a label of `spk_1`.

If an utterance from one speaker overlaps with an utterance from another speaker, Amazon Transcribe Medical orders them in the transcription by their start times. Utterances that overlap in the input audio don't overlap in the transcription output.

You can enable speaker diarization when you transcribe an audio file using batch transcription job, or in a real-time stream.

Topics

- [Identifying speakers and labeling their speech in audio files \(p. 296\)](#)
- [Identifying speakers and labeling their speech in real-time streams \(p. 299\)](#)

Identifying speakers and labeling their speech in audio files

You can enable speaker identification in a batch transcription job using either the [StartMedicalTranscriptionJob](#) API or the AWS Management Console. This enables you to identify the speakers in a clinician-patient conversation and determine who said what in the transcription output.

AWS Management Console

To use the AWS Management Console to enable speaker diarization in your transcription job, you enable audio identification and then speaker identification.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, provide information about your transcription job.
5. Choose **Next**.
6. Enable **Audio identification**.
7. For **Audio identification type**, choose **Speaker identification**.
8. For **Maximum number of speakers**, enter the maximum number of speakers that you think are speaking in your audio file. For best results, match the number of speakers that you ask to identify to the number of speakers in the input audio. If you specify a value less than the number of speakers in your input audio, the transcription text of the most similar sounding speakers are attributed to a speaker label.
9. Choose **Create**.

API

To identify speakers in an audio file using a batch transcription job (API)

- For the `StartMedicalTranscriptionJob` API, specify the following.
 - a. For `MedicalTranscriptionJobName`, specify a name that is unique in your AWS account.
 - b. For `LanguageCode`, specify the language code that corresponds to the language spoken in the audio file.
 - c. For the `MediaFileUri` parameter of the `Media` object, specify the name of the audio file that you want to transcribe.
 - d. For `Specialty`, specify the medical specialty of the clinician speaking in the audio file.
 - e. For `Type`, specify `CONVERSATION`.
 - f. For `OutputBucketName`, specify the Amazon S3 bucket to store the transcription results.
 - g. For the `Settings` object, specify the following.
 - i. `ShowSpeakerLabels` – `true`.
 - ii. `MaxSpeakerLabels` – An integer between 2 and 10 to indicate the number of speakers that you think are speaking in your audio. For best results, match the number of speakers that you ask to identify to the number of speakers in the input audio. If you specify a value less than the number of speakers in your input audio, the transcription text of the most similar sounding speakers are attributed to the same speaker label.

The following request uses the AWS SDK for Python (Boto3) to start a batch transcription job of a primary care clinician patient dialogue with speaker identification enabled.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
transcribe.start_medical_transcription_job(
    MedicalTranscriptionJobName = job_name,
    Media={
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'my-output-files/',
    LanguageCode = 'en-US',
    Specialty = 'PRIMARYCARE',
    Type = 'CONVERSATION',
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    Settings = {'ShowSpeakerLabels': True,
               'MaxSpeakerLabels': 2
               }
)
while True:
    status = transcribe.get_medical_transcription_job(MedicalTranscriptionJobName =
    job_name)
    if status['MedicalTranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED',
    'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

The following example code shows the transcription results of a transcription job with speaker identification enabled.

```
{
  "jobName": "job ID",
  "accountId": "111122223333",
  "results": {
    "transcripts": [
      {
        "transcript": "Professional answer."
      }
    ],
    "speaker_labels": {
      "speakers": 1,
      "segments": [
        {
          "start_time": "0.000000",
          "speaker_label": "spk_0",
          "end_time": "1.430",
          "items": [
            {
              "start_time": "0.100",
              "speaker_label": "spk_0",
              "end_time": "0.690"
            },
            {
              "start_time": "0.690",
              "speaker_label": "spk_0",
              "end_time": "1.210"
            }
          ]
        }
      ]
    },
    "items": [
      {
        "start_time": "0.100",
        "end_time": "0.690",
        "alternatives": [
          {
            "confidence": "0.8162",
            "content": "Professional"
          }
        ],
        "type": "pronunciation"
      },
      {
        "start_time": "0.690",
        "end_time": "1.210",
        "alternatives": [
          {
            "confidence": "0.9939",
            "content": "answer"
          }
        ],
        "type": "pronunciation"
      },
      {
        "alternatives": [
          {
            "content": "."
          }
        ],
        "type": "punctuation"
      }
    ]
  }
}
```



```
    }  
  ],  
  "status": "COMPLETED"  
}
```

AWS CLI

To transcribe an audio file of a conversation between a clinician practicing primary care and a patient in an audio file and identify what each person said in the transcription output (AWS CLI)

- Run the following code.

```
aws transcribe start-transcription-job \  
--region us-west-2 \  
--cli-input-json file://example-start-command.json
```

The following code shows the contents of `example-start-command.json`.

```
{  
  "MedicalTranscriptionJobName": "my-first-med-transcription-job",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac"  
  },  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "OutputKey": "my-output-files/",  
  "LanguageCode": "en-US",  
  "Specialty": "PRIMARYCARE",  
  "Type": "CONVERSATION",  
  "Settings": {  
    "ShowSpeakerLabels": true,  
    "MaxSpeakerLabels": 2  
  }  
}
```

Identifying speakers and labeling their speech in real-time streams

To identify speakers and label their speech in a real-time stream, use the AWS Management Console or a streaming request. Speaker identification works best for identifying between two and five speakers in a stream. Although Amazon Transcribe Medical can identify more than five speakers in a stream, the accuracy of speaker identification decreases if you exceed that number.

To start an HTTP/2 request, use the [StartMedicalStreamTranscription](#) API. To start a WebSocket request, use a pre-signed URI. The URI contains the information required to set up bi-directional communication between your application and Amazon Transcribe Medical.

Identifying speakers in audio that is spoken into your microphone (AWS Management Console)

You can use the AWS Management Console to start a real-time stream of a clinician-patient conversation, or a dictation that is spoken into your microphone in real-time.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, for Amazon Transcribe Medical choose **Real-time transcription**.
3. For **Audio input type**, choose the type of medical speech that you want to transcribe.
4. For **Additional settings**, choose **Speaker identification**.
5. Choose **Start streaming** to start transcribing your real-time audio.
6. Speak into the microphone.

Identifying speakers in an HTTP/2 stream

To identify speakers in an HTTP/2 stream of a medical conversation, use the [StartMedicalStreamTranscription](#) API and specify the following:

- For `LanguageCode`, specify the language code that corresponds to the language in the stream. The valid value is `en-US`.
- For `MediaSampleHertz`, specify the sample rate of the audio.
- For `Specialty`, specify the medical specialty of the provider.
- `ShowSpeakerLabel` – `true`

For more information on setting up an HTTP/2 stream to transcribe a medical conversation, see [Streaming request \(p. 270\)](#).

Identifying speakers in a WebSocket request

To identify speakers in WebSocket streams with the API, use the following format to create a pre-signed URI to start a WebSocket request and set `show-speaker-label` to `true`.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/medical-stream-transcription-  
websocket  
?language-code=languageCode  
&X-Amz-Algorithm=AWS4-HMAC-SHA256  
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request  
&X-Amz-Date=20220208T235959Z  
&X-Amz-Expires=250  
&X-Amz-Security-Token=security-token  
&X-Amz-Signature=Signature Version 4 signature  
&X-Amz-SignedHeaders=host  
&media-encoding=flac  
&sample-rate=16000  
&session-id=sessionId  
&specialty=medicalSpecialty  
&type=CONVERSATION  
&vocabulary-name=vocabularyName  
&show-speaker-label=boolean
```

The following code shows the truncated example response of a streaming request.

```
{  
  "Transcript": {  
    "Results": [  
      {  
        "Alternatives": [  
          {  
            "Items": [  
              {  
                "Confidence": 0.97,  

```

```
    "Content": "From",
    "EndTime": 18.98,
    "Speaker": "0",
    "StartTime": 18.74,
    "Type": "pronunciation",
    "VocabularyFilterMatch": false
  },
  {
    "Confidence": 1,
    "Content": "the",
    "EndTime": 19.31,
    "Speaker": "0",
    "StartTime": 19,
    "Type": "pronunciation",
    "VocabularyFilterMatch": false
  },
  {
    "Confidence": 1,
    "Content": "last",
    "EndTime": 19.86,
    "Speaker": "0",
    "StartTime": 19.32,
    "Type": "pronunciation",
    "VocabularyFilterMatch": false
  },
  ...
  {
    "Confidence": 1,
    "Content": "chronic",
    "EndTime": 22.55,
    "Speaker": "0",
    "StartTime": 21.97,
    "Type": "pronunciation",
    "VocabularyFilterMatch": false
  },
  ...
  {
    "Confidence": 1,
    "Content": "fatigue",
    "EndTime": 24.42,
    "Speaker": "0",
    "StartTime": 23.95,
    "Type": "pronunciation",
    "VocabularyFilterMatch": false
  },
  {
    "EndTime": 25.22,
    "StartTime": 25.22,
    "Type": "speaker-change",
    "VocabularyFilterMatch": false
  },
  {
    "Confidence": 0.99,
    "Content": "True",
    "EndTime": 25.63,
    "Speaker": "1",
    "StartTime": 25.22,
    "Type": "pronunciation",
    "VocabularyFilterMatch": false
  },
  {
    "Content": ".",
    "EndTime": 25.63,
    "StartTime": 25.63,
    "Type": "punctuation",
    "VocabularyFilterMatch": false
  }
}
```

```
        ],
        "Transcript": "From the last note she still has mild sleep deprivation and
chronic fatigue True."
    }
  ],
  "EndTime": 25.63,
  "IsPartial": false,
  "ResultId": "XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX",
  "StartTime": 18.74
}
]
}
```

Amazon Transcribe Medical breaks your incoming audio stream based on natural speech segments, such as a change in speaker or a pause in the audio. The transcription is returned progressively to your application, with each response containing more transcribed speech until the entire segment is transcribed. The preceding code is a truncated example of a fully-transcribed speech segment. Speaker labels only appear for entirely transcribed segments.

The following list shows the organization of the objects and parameters in a streaming transcription output.

Transcript

Each speech segment has its own `Transcript` object.

Results

Each `Transcript` object has its own `Results` object. This object contains the `isPartial` field. When its value is `false`, the results returned are for an entire speech segment.

Alternatives

Each `Results` object has an `Alternatives` object.

Items

Each `Alternatives` object has its own `Items` object that contains information about each word and punctuation mark in the transcription output. When you enable speaker identification, each word has a `Speaker` label for fully-transcribed speech segments. Amazon Transcribe Medical uses this label to assign a unique integer to each speaker it identifies in the stream. The `Type` parameter having a value of `speaker-change` indicates that one person has stopped speaking and that another person is about to begin.

Transcript

Each `Items` object contains a transcribed speech segment as the value of the `Transcript` field.

For more information about WebSocket requests, see [Creating a pre-signed URI \(p. 284\)](#).

Transcribing multi-channel audio

If you have an audio file or stream that has multiple channels, you can use *channel identification* to transcribe the speech from each of those channels. Amazon Transcribe Medical transcribes the speech from each channel separately. It combines the separate transcriptions of each channel into a single transcription output.

Use channel identification to identify the separate channels in your audio and transcribe the speech from each of those channels. Enable this in situations such as a caller and agent scenario. Use this to distinguish a caller from an agent in recordings or streams from contact centers that perform drug safety monitoring.

You can enable channel identification for both batch processing and real-time streaming. The following list describes how to enable it for each method.

- Batch transcription – AWS Management Console and [StartMedicalTranscriptionJob](#) API
- Streaming transcription – WebSocket streaming and [StartMedicalStreamTranscription](#) API

Transcribing multi-channel audio files

When you transcribe an audio file, Amazon Transcribe Medical returns a list of *items* for each channel. An item is a transcribed word or punctuation mark. Each word has a start time and an end time. If a person on one channel speaks over a person on a separate channel, the start times and end times of the items for each channel overlap while the individuals are speaking over each other.

By default, you can transcribe audio files with two channels. You can request a quota increase if you need to transcribe files that have more than two channels. For information about requesting a quota increase, see [AWS service quotas](#).

To transcribe multi-channel audio in a batch transcription job, use the AWS Management Console or the [StartMedicalTranscriptionJob](#) API.

AWS Management Console

To use the AWS Management Console to enable channel identification in your batch transcription job, you enable audio identification and then channel identification. Channel identification is a subset of audio identification in the AWS Management Console.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, provide information about your transcription job.
5. Choose **Next**.
6. Enable **Audio identification**.
7. For **Audio identification type**, choose **Channel identification**.
8. Choose **Create**.

API

To transcribe a multi-channel audio file (API)

- For the [StartMedicalTranscriptionJob](#) API, specify the following.
 - a. For `TranscriptionJobName`, specify a name unique to your AWS account.
 - b. For `LanguageCode`, specify the language code that corresponds to the language spoken in the audio file. The valid value is `en-US`.
 - c. For the `MediaFileUri` parameter of the `Media` object, specify the name of the media file that you want to transcribe.
 - d. For the `Settings` object, set `ChannelIdentification` to `true`.

The following is an example request using the AWS SDK for Python (Boto3).

```
from __future__ import print_function
import time
import boto3
```

```
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-transcription-job"
job_name = "my-first-med-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
transcribe.start_medical_transcription_job(
    MedicalTranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'output-files/',
    LanguageCode = 'en-US',
    Specialty = 'PRIMARYCARE',
    Type = 'CONVERSATION',
    Settings = {
        'ChannelIdentification': True
    }
)
while True:
    status = transcribe.get_transcription_job(MedicalTranscriptionJobName = job_name)
    if status['MedicalTranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED',
        'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

AWS CLI

To transcribe a multi-channel audio file using a batch transcription job (AWS CLI)

- Run the following code.

```
aws transcribe start-medical-transcription-job \
--region us-west-2 \
--cli-input-json file://example-start-command.json
```

The following is the code of `example-start-command.json`.

```
{
  "MedicalTranscriptionJobName": "my-first-med-transcription-job",
  "Media": {
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac"
  },
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",
  "OutputKey": "my-output-files/",
  "LanguageCode": "en-US",
  "Specialty": "PRIMARYCARE",
  "Type": "CONVERSATION",

  "Settings":{
    "ChannelIdentification": true
  }
}
```

The following code shows the transcription output for an audio file that has a conversation on two channels.

```
{
  "jobName": "job id",
  "accountId": "111122223333",
  "results": {
    "transcripts": [
      {
        "transcript": "When you try ... It seems to ..."
      }
    ],
    "channel_labels": {
      "channels": [
        {
          "channel_label": "ch_0",
          "items": [
            {
              "start_time": "12.282",
              "end_time": "12.592",
              "alternatives": [
                {
                  "confidence": "1.0000",
                  "content": "When"
                }
              ],
              "type": "pronunciation"
            },
            {
              "start_time": "12.592",
              "end_time": "12.692",
              "alternatives": [
                {
                  "confidence": "0.8787",
                  "content": "you"
                }
              ],
              "type": "pronunciation"
            },
            {
              "start_time": "12.702",
              "end_time": "13.252",
              "alternatives": [
                {
                  "confidence": "0.8318",
                  "content": "try"
                }
              ],
              "type": "pronunciation"
            },
            ...
          ]
        },
        {
          "channel_label": "ch_1",
          "items": [
            {
              "start_time": "12.379",
              "end_time": "12.589",
              "alternatives": [
                {
                  "confidence": "0.5645",
                  "content": "It"
                }
              ],
              "type": "pronunciation"
            },
          ],
        }
      ]
    }
  }
}
```

```
{
  "start_time": "12.599",
  "end_time": "12.659",
  "alternatives": [
    {
      "confidence": "0.2907",
      "content": "seems"
    }
  ],
  "type": "pronunciation"
},
{
  "start_time": "12.669",
  "end_time": "13.029",
  "alternatives": [
    {
      "confidence": "0.2497",
      "content": "to"
    }
  ],
  "type": "pronunciation"
},
  ...
]
}
```

Transcribing multi-channel audio streams

You can transcribe audio from separate channels in either HTTP/2 or WebSocket streams using the [StartMedicalStreamTranscription](#) API.

By default, you can transcribe streams with two channels. You can request a quota increase if you need to transcribe streams that have more than two channels. For information about requesting a quota increase, see [AWS service quotas](#).

Transcribing multi-channel audio in an HTTP/2 stream

To transcribe multi-channel audio in an HTTP/2 stream, use the [StartMedicalStreamTranscription](#) API and specify the following:

- `LanguageCode` – The language code of the audio. The valid value is `en-US`.
- `MediaEncoding` – The encoding of the audio. Valid values are `ogg-opus`, `flac`, and `pcm`.
- `EnableChannelIdentification` – `true`
- `NumberOfChannels` – the number of channels in your streaming audio.

For more information on setting up an HTTP/2 stream to transcribe a medical conversation, see [Streaming request \(p. 270\)](#).

Transcribing multi-channel audio in a WebSocket stream

To identify speakers in WebSocket streams, use the following format to create a pre-signed URI and start a WebSocket request. Specify `enable-channel-identification` as `true` and the number of channels in your stream in `number-of-channels`. A pre-signed URI contains the information needed to set up bi-directional communication between your application and Amazon Transcribe Medical.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/medical-stream-transcription-  
websocket  
?language-code=languageCode
```



```
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=Signature Version 4 signature
&X-Amz-SignedHeaders=host
&media-encoding=flac
&sample-rate=16000
&session-id=sessionId
&enable-channel-identification=true
&number-of-channels=2
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

For more information about WebSocket requests, see [Creating a pre-signed URI \(p. 284\)](#).

Multi-channel streaming output

The output of a streaming transcription is the same for HTTP/2 and WebSocket requests. The following is an example output.

```
{
  "resultId": "XXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXX",
  "startTime": 0.11,
  "endTime": 0.66,
  "isPartial": false,
  "alternatives": [
    {
      "transcript": "Left.",
      "items": [
        {
          "startTime": 0.11,
          "endTime": 0.45,
          "type": "pronunciation",
          "content": "Left",
          "vocabularyFilterMatch": false
        },
        {
          "startTime": 0.45,
          "endTime": 0.45,
          "type": "punctuation",
          "content": ".",
          "vocabularyFilterMatch": false
        }
      ]
    }
  ],
  "channelId": "ch_0"
}
```

For each speech segment, there is a `channelId` flag that indicates which channel the speech belongs to.

Transcribing a medical dictation

You can use Amazon Transcribe Medical to transcribe clinician-dictated medical notes using either a batch transcription job or a real-time stream. Batch transcription jobs enable you to transcribe audio files. You specify the medical specialty of the clinician in your transcription job or stream to ensure that Amazon Transcribe Medical produces transcription results with the highest possible accuracy.

You can transcribe a medical dictation in the following specialties:

- Cardiology – available in streaming transcription only
- Neurology – available in streaming transcription only
- Oncology – available in streaming transcription only
- Primary Care – includes the following types of medical practice:
 - Family medicine
 - Internal medicine
 - Obstetrics and Gynecology (OB-GYN)
 - Pediatrics
- Radiology – available in streaming transcription only
- Urology – available in streaming transcription only

You can improve transcription accuracy by using custom vocabularies. For information on how medical custom vocabularies work, see [Improving transcription accuracy with medical custom vocabularies \(p. 313\)](#).

By default, Amazon Transcribe Medical returns the transcription with the highest confidence level. If you'd like to configure it to return alternative transcriptions, see [Generating alternative transcriptions \(p. 327\)](#).

For information about how numbers and medical measurements appear in the transcription output, see [Transcribing numbers \(p. 259\)](#) and [Transcribing medical terms and measurements \(p. 260\)](#).

Topics

- [Transcribing an audio file of a medical dictation \(p. 308\)](#)
- [Transcribing a medical dictation in a real-time stream \(p. 311\)](#)

Transcribing an audio file of a medical dictation

Use a batch transcription job to transcribe audio files of medical conversations. You can use this to transcribe a clinician-patient dialogue. You can start a batch transcription job in either the [StartMedicalTranscriptionJob](#) API or the AWS Management Console.

When you start a medical transcription job with the [StartMedicalTranscriptionJob](#) API, you specify `PRIMARYCARE` as the value of the `Specialty` parameter.

AWS Management Console

To transcribe a clinician-patient dialogue (AWS Management Console)

To use the AWS Management Console to transcribe a clinician-patient dialogue, create a transcription job and choose **Conversation** for **Audio input type**.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, under **Job settings**, specify the following.
 - a. **Name** – the name of the transcription job.
 - b. **Audio input type** – **Dictation**
5. For the remaining fields, specify the Amazon S3 location of your audio file and where you want to store the output of your transcription job.

6. Choose **Next**.
7. Choose **Create**.

API

To transcribe a medical conversation using a batch transcription job (API)

- For the [StartMedicalTranscriptionJob](#) API, specify the following.
 - a. For `MedicalTranscriptionJobName`, specify a name unique in your AWS account.
 - b. For `LanguageCode`, specify the language code that corresponds to the language spoken in your audio file and the language of your vocabulary filter.
 - c. In the `MediaFileUri` parameter of the `Media` object, specify the name of the audio file that you want to transcribe.
 - d. For `Specialty`, specify the medical specialty of the clinician speaking in the audio file.
 - e. For `Type`, specify `DICTATION`.
 - f. For `OutputBucketName`, specify the Amazon S3 bucket to store the transcription results.

The following is an example request that uses the AWS SDK for Python (Boto3) to transcribe a medical dictation of a clinician in the `PRIMARYCARE` specialty.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe')
job_name = "my-first-med-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac"
transcribe.start_medical_transcription_job(
    MedicalTranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'my-output-files/',
    LanguageCode = 'en-US',
    Specialty = 'PRIMARYCARE',
    Type = 'DICTATION'
)
while True:
    status = transcribe.get_medical_transcription_job(MedicalTranscriptionJobName =
    job_name)
    if status['MedicalTranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED',
    'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

The following example code shows the transcription results of a medical dictation.

```
{
  "jobName": "dictation-medical-transcription-job",
  "accountId": "111122223333",
  "results": {
```

```
"transcripts": [
  {
    "transcript": "... came for a follow up visit today..."
  }
],
"items": [
  {
    ...
    "start_time": "4.85",
    "end_time": "5.12",
    "alternatives": [
      {
        "confidence": "1.0",
        "content": "came"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "5.12",
    "end_time": "5.29",
    "alternatives": [
      {
        "confidence": "1.0",
        "content": "for"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "5.29",
    "end_time": "5.33",
    "alternatives": [
      {
        "confidence": "0.9955",
        "content": "a"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "5.33",
    "end_time": "5.66",
    "alternatives": [
      {
        "confidence": "0.9754",
        "content": "follow"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "5.66",
    "end_time": "5.75",
    "alternatives": [
      {
        "confidence": "0.9754",
        "content": "up"
      }
    ],
    "type": "pronunciation"
  },
  {
    "start_time": "5.75",
    "end_time": "6.02",
    "alternatives": [
```

```
        {
            "confidence": "1.0",
            "content": "visit"
        }
    ]
    ...
},
"status": "COMPLETED"
}
```

AWS CLI

To identify the speakers in an audio file using a batch transcription job (AWS CLI)

- Run the following code.

```
aws transcribe start-medical-transcription-job \
--region us-west-2 \
--cli-input-json file://example-start-command.json
```

The following code shows the contents of `example-start-command.json`.

```
{
  "MedicalTranscriptionJobName": "my-first-med-transcription-job",
  "Media": {
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac"
  },
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",
  "OutputKey": "my-output-files/",
  "LanguageCode": "en-US",
  "Specialty": "PRIMARYCARE",
  "Type": "DICTATION"
}
```

Transcribing a medical dictation in a real-time stream

Use a WebSocket stream to transcribe a medical dictation as an audio stream. You can also use the AWS Management Console to transcribe speech that you or others speak directly into a microphone.

For an HTTP/2 or a WebSocket stream, you can transcribe audio in the following medical specialties:

- Cardiology
- Oncology
- Neurology
- Primary Care
- Radiology
- Urology

Each medical specialty includes many types of procedures and appointments. Clinicians therefore dictate many different types of notes. Use the following examples as guidance to help you specify the value of the `specialty` URI parameter of the WebSocket request, or the `Specialty` parameter of the [StartMedicalStreamTranscription](#) API:

- For a dictation after electrophysiology or echocardiogram procedure, choose `CARDIOLOGY`.
- For a dictation after a surgical oncology or radiation oncology procedure, choose `ONCOLOGY`.
- For a physician dictating notes indicating a diagnosis of encephalitis, choose `NEUROLOGY`.
- For a dictation of procedure notes to break up a bladder stone, choose `UROLOGY`.
- For a dictation of clinician notes after an internal medicine consultation, choose `PRIMARYCARE`.
- For a dictation of a physician communicating the findings of a CT scan, PET scan, MRI, or radiograph, choose `RADIOLOGY`.
- For a dictation of physician notes after a gynecology consultation, choose `PRIMARYCARE`.

To improve transcription accuracy of specific terms in a real-time stream, use a custom vocabulary. To enable a custom vocabulary, set the value of `vocabulary-name` to the name of the custom vocabulary you want to use.

Transcribing a dictation spoken into your microphone with the AWS Management Console

To use the AWS Management Console to transcribe streaming audio of a medical dictation, choose the option to transcribe a medical dictation, start the stream, and begin speaking into the microphone.

To transcribe streaming audio of a medical dictation (AWS Management Console)

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Real-time transcription**.
3. Choose **Dictation**.
4. For **Medical specialty**, choose the medical specialty of the clinician speaking in the stream.
5. Choose **Start streaming**.
6. Speak into the microphone.

Transcribing a dictation in an HTTP/2 stream

To transcribe an HTTP/2 stream of a medical dictation, use the [StartMedicalStreamTranscription](#) API and specify the following:

- `LanguageCode` – The language code. The valid value is `en-US`
- `MediaEncoding` – The encoding used for the input audio. Valid values are `pcm`, `ogg-opus`, and `flac`.
- `Specialty` – The specialty of the medical professional.
- `Type` – `DICTATION`

For more information on setting up an HTTP/2 stream to transcribe a medical dictation, see [Streaming request \(p. 270\)](#).

Using a WebSocket streaming request to transcribe a medical dictation

To transcribe a medical dictation in a real-time stream using a WebSocket request, you create a pre-signed URI. This URI contains the information needed to set up the audio stream between your application and Amazon Transcribe Medical. For more information on creating WebSocket requests, see [Establish a bi-directional connection using the WebSocket protocol \(p. 283\)](#).

Use the following template to create your pre-signed URI.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/medical-stream-transcription-  
websocket  
?language-code=languageCode
```

```
&X-Amz-Algorithm=AWS4-HMAC-SHA256
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request
&X-Amz-Date=20220208T235959Z
&X-Amz-Expires=250
&X-Amz-Security-Token=security-token
&X-Amz-Signature=Signature Version 4 signature
&X-Amz-SignedHeaders=host
&media-encoding=flac
&sample-rate=16000
&session-id=sessionId
&specialty=medicalSpecialty
&type=DICTATION
&vocabulary-name=vocabularyName
&show-speaker-label=boolean
```

For more information on creating pre-signed URIs, see [Creating a pre-signed URI \(p. 284\)](#).

Improving transcription accuracy with medical custom vocabularies

To improve transcription accuracy in Amazon Transcribe Medical, create and use one or more medical custom vocabularies. A *custom vocabulary* is a collection of words or phrases that are domain-specific. This collection helps improve the performance of Amazon Transcribe Medical in transcribing those words or phrases.

You are responsible for the integrity of your own data when you use Amazon Transcribe Medical. Do not enter confidential information, personal information (PII), or protected health information (PHI), into a custom vocabulary.

For best results, create separate small custom vocabularies that each help transcribe a specific audio recording. You receive greater improvements in transcription accuracy than if you created one large custom vocabulary to use with all of your recordings.

By default, you can have up to 100 custom vocabularies in your AWS account. A custom vocabulary can't exceed 50 KB in size. For information on requesting an increase to the number of custom vocabularies that you can have in your AWS account, see [AWS service quotas](#).

Custom vocabularies are available in US English (en-US).

Topics

- [Creating a text file for your medical custom vocabulary \(p. 313\)](#)
- [Using a text file to create a medical custom vocabulary \(p. 316\)](#)
- [Transcribing an audio file using a medical custom vocabulary \(p. 318\)](#)
- [Transcribing a real-time stream using a medical custom vocabulary \(p. 319\)](#)
- [Character sets for Amazon Transcribe Medical \(p. 321\)](#)

Creating a text file for your medical custom vocabulary

To create a custom vocabulary, you create a text file that is in UTF-8 format. In this file, you create a four column table, with each column specifying a field. Each field tells Amazon Transcribe Medical either how the domain-specific terms are pronounced or how to display these terms in your transcriptions. You store the text file containing these fields in an Amazon S3 bucket.

Understanding how to format your text file

To create a medical custom vocabulary, you enter the column names as a header row. You enter the values for each column beneath the header row.

The following are the names of the four columns of the table:

- `Phrase` – column required, values required
- `IPA` – column required, values can be optional
- `SoundsLike` – column required, values can be optional
- `DisplayAs` – column required, values can be optional

When you create a custom vocabulary, make sure that you:

- Separate each column with a single Tab character. Amazon Transcribe throws an error message if you try to separate the columns with spaces or multiple Tab characters.
- Make sure that there's no trailing spaces or white space after each value within a column.

Make sure that the values that you enter for each column:

- Have fewer than 256 characters, including hyphens
- Use only characters from the allowed character set, see [Character sets for Amazon Transcribe Medical \(p. 321\)](#).

Entering values for the columns of the table

The following information shows you how to specify values for the four columns of the table:

- **Phrase** – The word or phrase that should be recognized. You must enter values in this column.

If the entry is a phrase, separate the words with a hyphen (-). For example, enter **cerebral autosomal dominant arteriopathy with subcortical infarcts and leukoencephalopathy** as **cerebral-autosomal-dominant-arteriopathy-with-subcortical-infarcts-and-leukoencephalopathy**.

Enter acronyms or other words whose letters should be pronounced individually as single letters followed by dots, such as **D.N.A.** or **S.T.E.M.I.**. To enter the plural form of an acronym, such as "STEMIs," separate the "s" from the acronym with a hyphen: "**S.T.E.M.I-s**" You can use either uppercase or lowercase letters for acronyms.

The `Phrase` column is required. You can use any of the allowed characters for the input language. For allowed characters, see [Character sets for Amazon Transcribe Medical \(p. 321\)](#). If you don't specify the `DisplayAs` column, Amazon Transcribe Medical uses the contents of the `Phrase` column in the output file.

- **IPA** (column required, values can be optional) – To specify the pronunciation of a word or phrase, you can include characters in the [International Phonetic Alphabet \(IPA\)](#) in this column. The `IPA` column can't contain leading or trailing spaces, and you must use a single space to separate each phoneme in the input. For example, in English you would enter the phrase **acute-respiratory-distress-syndrome** as **# k j u t # # s p # # # t # # i d # s t # # s s # n d # o # m**. You would enter the phrase **A.L.I.** as **e# # l # l**.

Even if you don't specify the contents of the `IPA` column, you must include a blank `IPA` column. If you include values in the `IPA` column, you can't provide values for the `SoundsLike` column.

For a list of allowed IPA characters for a specific language, see [Character sets for Amazon Transcribe Medical \(p. 321\)](#). US English is the only language available in Amazon Transcribe Medical.

- **SoundsLike** (column required, values can be optional) – You can break a word or phrase down into smaller segments and provide a pronunciation for each segment using the standard orthography of the language to mimic the way that the word sounds. For example, you can provide pronunciation hints for the phrase **cerebral-autosomal-dominant-arteriopathy-with-subcortical-infarcts-and-leukoencephalopathy** like this: **sir-e-brul-aut-o-som-ul-dah-mi-nant-ar-ter-ri-o-pa-thy-with-sub-cor-ti-cul-in-farcts-and-lewk-o-en-ce-phul-ah-pu-thy**. The hint for the phrase **atrioventricular-nodal-reentrant-tachycardia** would look like this: **ay-tree-o-ven-trick-u-lar-node-al-re-entr-ant-tack-ih-card-ia**. You separate each part of the hint with a hyphen (-).

Even if you don't provide values for the SoundsLike column, you must include a blank SoundsLike column. If you include values in the SoundsLike column, you can't provide values for the IPA column.

You can use any of the allowed characters for the input language. For the list of allowed characters, see [Character sets for Amazon Transcribe Medical \(p. 321\)](#).

- **DisplayAs** (column required, values can be optional) – Defines how the word or phrase looks when it's output. For example, if the word or phrase is **cerebral-autosomal-dominant-arteriopathy-with-subcortical-infarcts-and-leukoencephalopathy**, you can specify the display form as cerebral autosomal dominant arteriopathy with subcortical infarcts and leukoencephalopathy, so that the hyphen is not present. You can also specify DisplayAs as CADASIL if you'd like to show the acronym instead of the full term in the output.

If you don't specify the DisplayAs column, Amazon Transcribe Medical uses the Phrase column from the input file in the output.

You can use any UTF-8 character in the DisplayAs column.

You can include spaces only for the values in the IPA and DisplayAs columns.

To create the text file of your custom vocabulary, place each word or phrase in your text file on a separate line. Separate the columns with Tab characters. Include spaces only for values in the IPA and DisplayAs columns. Save the file with the extension `.txt` in an Amazon S3 bucket in the same AWS Region where you use Amazon Transcribe Medical to create your custom vocabulary.

If you edit your text file in Windows, make sure that your file is in LF format and not in CRLF format. Otherwise, you will be unable to create your custom vocabulary. Some text editors enable you to change the formatting with Find and Replace commands.

The following examples show text that you can use to create custom vocabularies. To create a custom vocabulary from these examples, copy an example into a text editor, replace [TAB] with a Tab character, and upload the saved text file to Amazon S3.

```
Phrase[TAB]IPA[TAB]SoundsLike[TAB]DisplayAs
acute-respiratory-distress-syndrome[TAB][TAB][TAB]acute respiratory distress syndrome
A.L.L.[TAB]e# # 1 # 1[TAB][TAB]ALL
atrioventricular-nodal-reentrant-tachycardia[TAB][TAB]ay-tree-o-ven-trick-u-lar-node-al-re-entr-ant-tack-ih-card-ia[TAB]
```

You can enter columns in any order. The following examples show other valid structures for the custom vocabulary input file.

```
Phrase[TAB]SoundsLike[TAB]IPA[TAB]DisplayAs
acute-respiratory-distress-syndrome[TAB][TAB][TAB]acute respiratory distress syndrome
A.L.L.[TAB][TAB]e# # 1 # 1[TAB]ALL
```

```
atrioventricular-nodal-reentrant-tachycardia[TAB]ay-tree-o-ven-trick-u-lar-node-al-re-entr-ant-tack-ih-card-ia[TAB][TAB]
```

```
DisplayAs[TAB]SoundsLike[TAB]IPA[TAB]Phrase
acute respiratory distress syndrome[TAB][TAB][TAB]acute-respiratory-distress-syndrome
ALL[TAB][TAB]e# # l # l[TAB]A.L.L.
[TAB]ay-tree-o-ven-trick-u-lar-node-al-re-entr-ant-tack-ih-card-ia[TAB]
[TAB]atrioventricular-nodal-reentrant-tachycardia
```

For reading ease, the following tables show the preceding examples more clearly in html format. They are meant only to illustrate the examples.

Phrase	IPA	SoundsLike	DisplayAs
acute-respiratory-distress-syndrome			acute respiratory distress syndrome
A.L.L.	eɪ ɛ l ɛ l		ALL
atrioventricular-nodal-reentrant-tachycardia		ay-tree-o-ven-trick-u-lar-node-al-re-entr-ant-tack-ih-card-ia	

Phrase	SoundsLike	IPA	DisplayAs
acute-respiratory-distress-syndrome			acute respiratory distress syndrome
atrioventricular-nodal-reentrant-tachycardia	ay-tree-o-ven-trick-u-lar-node-al-re-entr-ant-tack-ih-card-ia		
A.L.L.		eɪ ɛ l ɛ l	ALL

DisplayAs	SoundsLike	IPA	Phrase
acute respiratory distress syndrome			acute-respiratory-distress-syndrome
ALL		eɪ ɛ l ɛ l	A.L.L.
	ay-tree-o-ven-trick-u-lar-node-al-re-entr-ant-tack-ih-card-ia		atrioventricular-nodal-reentrant-tachycardia

Using a text file to create a medical custom vocabulary

To create a custom vocabulary, you must have prepared a text file that contains a collection a words or phrases. Amazon Transcribe Medical uses this text file to create a custom vocabulary that you can use to

improve the transcription accuracy of those words or phrases. You can create a custom vocabulary using the [CreateMedicalVocabulary](#) API or the Amazon Transcribe Medical console.

AWS Management Console

To use the AWS Management Console to create a custom vocabulary, you provide the Amazon S3 URI of the text file containing your words or phrases.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Custom vocabulary**.
3. For **Name**, under **Vocabulary settings**, choose a name for your custom vocabulary.
4. Specify the location of your audio file or video file in Amazon S3:
 - For **Vocabulary input file location on S3** under **Vocabulary settings**, specify the Amazon S3 URI that identifies the text file you will use to create your custom vocabulary.
 - For **Vocabulary input file location in S3**, choose **Browse S3** to browse for the text file and choose it.
5. Choose **Create vocabulary**.

You can see the processing status of your custom vocabulary in the AWS Management Console.

API

To create a medical custom vocabulary (API)

- For the [StartTranscriptionJob](#) API, specify the following.
 - a. For `LanguageCode`, specify `en-US`.
 - b. For `VocabularyFileUri`, specify the Amazon S3 location of the text file that you use to define your custom vocabulary.
 - c. For `VocabularyName`, specify a name for your custom vocabulary. The name you specify must be unique within your AWS account.

To see the processing status of your custom vocabulary, use the [GetMedicalVocabulary](#) API.

The following is an example request using the AWS SDK for Python (Boto3) to create a custom vocabulary.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
vocab_name = "my-first-vocabulary"
response = transcribe.create_medical_vocabulary(
    VocabularyName = vocab_name,
    VocabularyFileUri = 's3://DOC-EXAMPLE-BUCKET/my-vocabularies/my-vocabulary-table.txt'
    LanguageCode = 'en-US',
)

while True:
    status = transcribe.get_medical_vocabulary(VocabularyName = vocab_name)
    if status['VocabularyState'] in ['READY', 'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

AWS CLI

To identify the speakers in an audio file using a batch transcription job (AWS CLI)

- Run the following code.

```
aws transcribe create-medical-vocabulary \  
--vocabulary-name my-first-vocabulary \  
--vocabulary-file-uri s3://DOC-EXAMPLE-BUCKET/my-vocabularies/my-vocabulary-file.txt \  
--language-code en-US
```

Transcribing an audio file using a medical custom vocabulary

Use the [StartMedicalTranscriptionJob](#) or the AWS Management Console to start a transcription job that uses a custom vocabulary to improve transcription accuracy.

AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, provide information about your transcription job.
5. Choose **Next**.
6. Under **Customization**, enable **Custom vocabulary**.
7. Under **Vocabulary selection**, choose a custom vocabulary.
8. Choose **Create**.

API

To identify speakers in an audio file using a batch transcription job (API)

- For the [StartMedicalTranscriptionJob](#) API, specify the following.
 - a. For `MedicalTranscriptionJobName`, specify a name that is unique in your AWS account.
 - b. For `LanguageCode`, specify the language code that corresponds to the language spoken in your audio file and the language of your vocabulary filter.
 - c. For the `MediaFileUri` parameter of the `Media` object, specify the name of the audio file that you want to transcribe.
 - d. For `Specialty`, specify the medical specialty of the clinician speaking in the audio file.
 - e. For `Type`, specify whether the audio file is a conversation or a dictation.
 - f. For `OutputBucketName`, specify the Amazon S3 bucket to store the transcription results.
 - g. For the `Settings` object, specify the following.
 - `VocabularyName` – the name of your custom vocabulary.

The following request uses the AWS SDK for Python (Boto3) to start a batch transcription job with a custom vocabulary.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-med-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-media-file.flac"
transcribe.start_medical_transcription_job(
    MedicalTranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'my-output-files/',
    LanguageCode = 'en-US',
    Specialty = 'PRIMARYCARE',
    Type = 'CONVERSATION',
    Settings = {
        'VocabularyName': 'example-med-custom-vocab'
    }
)

while True:
    status = transcribe.get_medical_transcription_job(MedicalTranscriptionJobName =
job_name)
    if status['MedicalTranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED',
'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

Transcribing a real-time stream using a medical custom vocabulary

To improve transcription accuracy in a real-time stream, you can use a custom vocabulary using either HTTP/2 or WebSocket streams. To start an HTTP/2 request, use the [StartMedicalStreamTranscription](#) API. You can use a custom vocabulary in real-time using either the AWS Management Console, the [StartMedicalStreamTranscription](#) API, or by using the WebSocket protocol.

Transcribing a dictation that is spoken into your Microphone (AWS Management Console)

To use the AWS Management Console to transcribe streaming audio of a medical dictation, choose the option to transcribe a medical dictation, start the stream, and begin speaking into the microphone.

To transcribe streaming audio of a medical dictation (AWS Management Console)

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Real-time transcription**.
3. For **Medical specialty**, choose the medical specialty of the clinician speaking in the stream.
4. For **Audio input type**, choose either **Conversation** or **Dictation**.
5. For **Additional settings**, choose **Custom vocabulary**.
 - For **Vocabulary selection**, choose the custom vocabulary.

6. Choose **Start streaming**.
7. Speak into the microphone.

Identifying speakers in an HTTP/2 stream

The following is the syntax for the parameters of an HTTP/2 request.

```
POST /medical-stream-transcription HTTP/2
host: transcribestreaming.us-west-2.amazonaws.com
authorization: Generated value
x-amz-target: com.amazonaws.transcribe.Transcribe.StartMedicalStreamTranscription
x-amz-content-sha256: STREAMING-MED-AWS4-HMAC-SHA256-EVENTS
x-amz-date: 20220208T235959Z
x-amzn-transcribe-session-id: my-first-http2-med-stream
x-amzn-transcribe-language-code: en-US
x-amzn-transcribe-media-encoding: flac
x-amzn-transcribe-sample-rate: 16000
x-amzn-transcribe-vocabulary-name: my-first-med-vocab
x-amzn-transcribe-specialty: PRIMARYCARE
x-amzn-transcribe-type: CONVERSATION
x-amzn-transcribe-show-speaker-label: true
Content-type: application/vnd.amazon.eventstream
transfer-encoding: chunked
```

Parameter descriptions:

- **host:** Update the AWS Region ('us-west-2' in the preceding example) with the AWS Region you are calling. For a list of valid AWS Regions, see [AWS Regions and Endpoints](#).
- **authorization:** This is a generated field. To learn more about creating a signature, see [Signing AWS requests with Signature Version 4](#).
- **x-amz-target:** Don't alter this field; use the content shown in the preceding example.
- **x-amz-content-sha256:** This is a generated field. To learn more about calculating a signature, see [Signing AWS requests with Signature Version 4](#).
- **x-amz-date:** The date and time the signature is created. The format is YYYYMMDDTHHMMSSZ, where YYYY=year, MM=month, DD=day, HH=hour, MM=minute, SS=seconds, and 'T' and 'Z' are fixed characters. For more information, refer to [Handling Dates in Signature Version 4](#).
- **x-amzn-transcribe-session-id:** The name for your streaming session.
- **x-amzn-transcribe-language-code:** The encoding used for your input audio. Refer to [StartMedicalStreamTranscription](#) or [Supported languages and language-specific features \(p. 3\)](#) for a list of valid values.
- **x-amzn-transcribe-media-encoding:** The encoding used for your input audio. Valid values are pcm, ogg-opus, and flac.
- **x-amzn-transcribe-sample-rate:** The sample rate of the input audio (in Hertz). Amazon Transcribe supports a range from 8,000 Hz to 48,000 Hz. Low-quality audio, such as telephone audio, is typically around 8,000 Hz. High-quality audio typically ranges from 16,000 Hz to 48,000 Hz. Note that the sample rate you specify **must** match that of your audio.
- **x-amzn-transcribe-vocabulary-name:** The name of the vocabulary you want to use with your transcription.
- **x-amzn-transcribe-specialty:** The medical specialty being transcribed.
- **x-amzn-transcribe-type:** Choose whether this is a dictation or a conversation.
- **x-amzn-transcribe-show-speaker-label:** To enable diarization, this value must be true.
- **content-type:** Don't alter this field; use the content shown in the preceding example.

Identifying speakers in a WebSocket request

To identify speakers in WebSocket streams with the API, use the following format to create a pre-signed URI to start a WebSocket request and set `vocabulary-name` to the name of the custom vocabulary.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/medical-stream-transcription-  
websocket  
?language-code=en-US  
&X-Amz-Algorithm=AWS4-HMAC-SHA256  
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request  
&X-Amz-Date=20220208T235959Z  
&X-Amz-Expires=250  
&X-Amz-Security-Token=security-token  
&X-Amz-Signature=Signature Version 4 signature  
&X-Amz-SignedHeaders=host  
&media-encoding=flac  
&sample-rate=16000  
&session-id=sessionId  
&specialty=medicalSpecialty  
&type=CONVERSATION  
&vocabulary-name=vocabularyName  
&show-speaker-label=boolean
```

Character sets for Amazon Transcribe Medical

To use custom vocabularies in Amazon Transcribe Medical, use the following character sets.

English character set

For English custom vocabularies, you can use the following characters in the `Phrase` and `SoundsLike` columns:

- a - z
- A - Z
- ' (apostrophe)
- - (hyphen)
- . (period)

You can use the following International Phonetic Alphabet (IPA) characters in the `IPA` column of the vocabulary input file.

Character	Code	Character	Code
au	0061 028A	w	0077
ai	0061 026A	z	007A
b	0062	æ	00E6
d	0064	ð	00F0
ei	0065 026A	ŋ	014B
f	0066	ɑ	0251
g	0067	ɔ	0254

Character	Code	Character	Code
h	0068	ɔɪ	0254 026A
i	0069	ə	0259
j	006A	ɛ	025B
k	006B	ɜ	025D
l	006C	g	0261
ɹ	006C 0329	ɪ	026A
m	006D	ɹ	0279
n	006E	ʃ	0283
ŋ	006E 0329	ʊ	028A
oʊ	006F 028A	ʌ	028C
p	0070	ɹ	028D
s	0073	ʒ	0292
t	0074	ɔʒ	02A4
u	0075	ʈʃ	02A7
v	0076	θ	03B8

Identifying personal health information (PHI) in a transcription

Use *Personal Health Information Identification* to label personal health information (PHI) in your transcription results. By reviewing labels, you can find PHI that could be used to identify a patient.

You can identify PHI using either a real-time stream or batch transcription job.

You can use your own post-processing to redact the PHI identified in the transcription output.

Use *Personal Health Information Identification* to identify the following types of PHI:

- Personal PHI:
 - Names – Full name or last name and initial
 - Gender
 - Age
 - Phone numbers
 - Dates (not including the year) that directly relate to the patient
 - Email addresses
- Geographic PHI:
 - Physical address
 - Zip code
 - Name of medical center or practice

- Account PHI:
 - Fax numbers
 - Social security numbers (SSNs)
 - Health insurance beneficiary numbers
 - Account numbers
 - Certificate or license numbers
 - Biometric identifiers
 - Voice prints
- Vehicle PHI:
 - Vehicle identification number (VIN)
 - License plate number
- Other PHI:
 - Web Uniform Resource Location (URL)
 - Internet Protocol (IP) address numbers

Amazon Transcribe Medical is a Health Insurance Portability and Accountability Act of 1996 (HIPAA) compliant service. For more information, see [Amazon Transcribe Medical \(p. 256\)](#). For information about identifying PHI in an audio file, see [Identifying PHI in an audio file \(p. 323\)](#). For information about identifying PHI in a stream, see [Identifying PHI in a real-time stream \(p. 326\)](#).

Topics

- [Identifying PHI in an audio file \(p. 323\)](#)
- [Identifying PHI in a real-time stream \(p. 326\)](#)

Identifying PHI in an audio file

Use a batch transcription job to transcribe audio files and identify the personal health information (PHI) within them. When you activate Personal Health Information (PHI) Identification, Amazon Transcribe Medical labels the PHI that it identified in the transcription results. For information about the PHI that Amazon Transcribe Medical can identify, see [Identifying personal health information \(PHI\) in a transcription \(p. 322\)](#).

You can start a batch transcription job using either the [StartMedicalTranscriptionJob](#) API or the AWS Management Console.

AWS Management Console

To use the AWS Management Console to transcribe a clinician-patient dialogue, create a transcription job and choose **Conversation** for **Audio input type**.

To transcribe an audio file and identify its PHI (AWS Management Console)

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, under **Job settings**, specify the following.
 - a. **Name** – The name of the transcription job that is unique to your AWS account.
 - b. **Audio input type** – **Conversation** or **Dictation**.
5. For the remaining fields, specify the Amazon S3 location of your audio file and where you want to store the output of your transcription job.

6. Choose **Next**.
7. Under **Audio settings**, choose **PHI Identification**.
8. Choose **Create**.

API

To transcribe an audio file and identify its PHI using a batch transcription job (API)

- For the [StartMedicalTranscriptionJob](#) API, specify the following.
 - a. For `MedicalTranscriptionJobName`, specify a name that is unique to your AWS account.
 - b. For `LanguageCode`, specify the language code that corresponds to the language spoken in your audio file.
 - c. For the `MediaFileUri` parameter of the `Media` object, specify the name of the audio file that you want to transcribe.
 - d. For `Specialty`, specify the medical specialty of the clinician speaking in the audio file as `PRIMARYCARE`.
 - e. For `Type`, specify either `CONVERSATION` or `DICTATION`.
 - f. For `OutputBucketName`, specify the Amazon S3 bucket where you want to store the transcription results.

The following is an example request that uses the AWS SDK for Python (Boto3) to transcribe an audio file and identify the PHI of a patient.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe')
job_name = "my-first-transcription-job"
job_uri = "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac"
transcribe.start_medical_transcription_job(
    MedicalTranscriptionJobName = job_name,
    Media = {'MediaFileUri': job_uri},
    LanguageCode = 'en-US',
    ContentIdentificationType = 'PHI',
    Specialty = 'PRIMARYCARE',
    Type = 'type', # Specify 'CONVERSATION' for a medical conversation. Specify
'DICTATION' for a medical dictation.
    OutputBucketName = 'DOC-EXAMPLE-BUCKET'
)
while True:
    status = transcribe.get_medical_transcription_job(MedicalTranscriptionJobName =
job_name)
    if status['MedicalTranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED',
'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

The following example code shows the transcription results with patient PHI identified.

```
{
```

```
"jobName": "my-medical-transcription-job-name",
"accountId": "111122223333",
"results": {
  "transcripts": [{
    "transcript": "The patient's name is Bertrand."
  }],
  "items": [{
    "start_time": "0.0",
    "end_time": "0.37",
    "alternatives": [{
      "confidence": "0.9993",
      "content": "The"
    }],
    "type": "pronunciation"
  }, {
    "start_time": "0.37",
    "end_time": "0.44",
    "alternatives": [{
      "confidence": "0.9981",
      "content": "patient's"
    }],
    "type": "pronunciation"
  }, {
    "start_time": "0.44",
    "end_time": "0.52",
    "alternatives": [{
      "confidence": "1.0",
      "content": "name"
    }],
    "type": "pronunciation"
  }, {
    "start_time": "0.52",
    "end_time": "0.92",
    "alternatives": [{
      "confidence": "1.0",
      "content": "is"
    }],
    "type": "pronunciation"
  }, {
    "start_time": "0.92",
    "end_time": "0.9989",
    "alternatives": [{
      "confidence": "1.0",
      "content": "Bertrand"
    }],
    "type": "pronunciation"
  }, {
    "alternatives": [{
      "confidence": "0.0",
      "content": "."
    }],
    "type": "punctuation"
  }],
  "entities": [{
    "content": "Bertrand",
    "category": "PHI*-Personal*",
    "startTime": 0.92,
    "endTime": 1.2,
    "confidence": 0.9989
  }],
  "status": "COMPLETED"
}
```

AWS CLI

To transcribe an audio file and identify PHI using a batch transcription job (AWS CLI)

- Run the following code.

```
aws transcribe start-medical-transcription-job \  
--medical-transcription-job-name my-medical-transcription-job-name \  
--language-code en-US \  
--media MediaFileUri="s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac" \  
--output-bucket-name DOC-EXAMPLE-BUCKET \  
--specialty PRIMARYCARE \  
--type type \ # Choose CONVERSATION to transcribe a medical conversation.  
Choose DICTATION to transcribe a medical dictation.  
--content-identification-type PHI
```

Identifying PHI in a real-time stream

You can identify Personally Identifiable Health Information (PHI) in either HTTP/2 or WebSocket streams. When you activate PHI Identification, Amazon Transcribe Medical labels the PHI that it identifies in the transcription results. For information about the PHI that Amazon Transcribe Medical can identify, see [Identifying personal health information \(PHI\) in a transcription \(p. 322\)](#).

Identifying PHI in a dictation that is spoken into your microphone

To use the AWS Management Console to transcribe the speech picked up by your microphone and identify any PHI, choose **Dictation** as the audio input type, start the stream, and begin speaking into the microphone on your computer.

To identify PHI in a dictation using the AWS Management Console

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, choose **Real-time transcription**.
3. For **Audio input type**, choose **Dictation**.
4. For **Additional settings**, choose **PHI identification**.
5. Choose **Start streaming** and speak into the microphone.
6. Choose **Stop streaming** to end the dictation.

Identifying PHI in an HTTP/2 stream

To start an HTTP/2 stream with PHI Identification activated, use the [StartMedicalStreamTranscription](#) API and specify the following:

- For `LanguageCode`, specify the language code for the language spoken in the stream. For US English, specify `en-US`.
- For `MediaSampleHertz`, specify the sample rate of the audio.
- For `content-identification-type`, specify `PHI`.

Identifying PHI in a WebSocket stream

To start a WebSocket stream with PHI Identification activated, use the following format to create a pre-signed URL.

```
GET wss://transcribestreaming.us-west-2.amazonaws.com:8443/medical-stream-transcription-  
websocket?  
&X-Amz-Algorithm=AWS4-HMAC-SHA256  
&X-Amz-Credential=AKIAIOSFODNN7EXAMPLE%2F20220208%2Fus-west-2%2Ftranscribe%2Faws4_request  
&X-Amz-Date=20220208T235959Z  
&X-Amz-Expires=250  
&X-Amz-Security-Token=security-token  
&X-Amz-Signature=Signature Version 4 signature  
&X-Amz-SignedHeaders=host  
&language-code=en-US  
&media-encoding=flac  
&sample-rate=16000  
&specialty=medical-specialty  
&content-identification-type=PHI
```

Parameter definitions can be found in the [API Reference](#); parameters common to all AWS API operations are listed in the [Common Parameters](#) section.

Generating alternative transcriptions

When you use Amazon Transcribe Medical, you get the transcription that has the highest confidence level. However, you can configure Amazon Transcribe Medical to return additional transcriptions with lower confidence levels.

Use alternative transcriptions to see different interpretations of the transcribed audio. For example, in an application that enables a person to review the transcription, you can present the alternative transcriptions for the person to choose from.

You can generate alternative transcriptions with the AWS Management Console or the [StartMedicalTranscriptionJob](#) API.

AWS Management Console

To use the AWS Management Console to generate alternative transcriptions, you enable alternative results when you configure your job.

1. Sign in to the [AWS Management Console](#).
2. In the navigation pane, under Amazon Transcribe Medical, choose **Transcription jobs**.
3. Choose **Create job**.
4. On the **Specify job details** page, provide information about your transcription job.
5. Choose **Next**.
6. Enable **Alternative results**.
7. For **Maximum alternatives**, enter an integer value between 2 and 10, for the maximum number of alternative transcriptions you want in the output.
8. Choose **Create**.

API

To identify speakers in an audio file using a batch transcription job (API)

- For the [StartMedicalTranscriptionJob](#) API, specify the following.
 - a. For `MedicalTranscriptionJobName`, specify a name that is unique in your AWS account.
 - b. For `LanguageCode`, specify the language code that corresponds to the language spoken in your audio file and the language of your vocabulary filter.
 - c. In the `MediaFileUri` parameter of the `Media` object, specify the location of the audio file you want to transcribe.
 - d. For `Specialty`, specify the medical specialty of the clinician speaking in the audio file.
 - e. For `Type`, specify whether you're transcribing a medical conversation or a dictation.
 - f. For `OutputBucketName`, specify the Amazon S3 bucket to store the transcription results.
 - g. For the `Settings` object, specify the following.
 - i. `ShowAlternatives` - true.
 - ii. `MaxAlternatives` - An integer between 2 and 10 to indicate the number of alternative transcriptions you want in the transcription output.

The following request uses the AWS SDK for Python (Boto3) to start a transcription job that generates up to two alternative transcriptions.

```
from __future__ import print_function
import time
import boto3
transcribe = boto3.client('transcribe', 'us-west-2')
job_name = "my-first-transcription-job"
job_uri = s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac
transcribe.start_medical_transcription_job(
    MedicalTranscriptionJobName = job_name,
    Media = {
        'MediaFileUri': job_uri
    },
    OutputBucketName = 'DOC-EXAMPLE-BUCKET',
    OutputKey = 'my-output-files/',
    LanguageCode = 'en-US',
    Specialty = 'PRIMARYCARE',
    Type = 'CONVERSATION',
    Settings = {
        'ShowAlternatives': True,
        'MaxAlternatives': 2
    }
)

while True:
    status = transcribe.get_medical_transcription_job(MedicalTranscriptionJobName =
job_name)
    if status['MedicalTranscriptionJob']['TranscriptionJobStatus'] in ['COMPLETED',
'FAILED']:
        break
    print("Not ready yet...")
    time.sleep(5)
print(status)
```

AWS CLI

To transcribe an audio file of a conversation between a primary care clinician and a patient in an audio file and identify what each person said in the transcription output (AWS CLI)

- Run the following code.

```
aws transcribe start-transcription-job \  
--cli-input-json file://filepath/example-start-command.json
```

The following code shows the contents of `example-start-command.json`.

```
>  
{  
  "MedicalTranscriptionJobName": "my-first-transcription-job",  
  "LanguageCode": "en-US",  
  "Specialty": "PRIMARYCARE",  
  "Type": "CONVERSATION",  
  "OutputBucketName": "DOC-EXAMPLE-BUCKET",  
  "Media": {  
    "MediaFileUri": "s3://DOC-EXAMPLE-BUCKET/my-input-files/my-audio-file.flac"  
  },  
  "Settings": {  
    "ShowAlternatives": true,  
    "MaxAlternatives": 2  
  }  
}
```

Guidelines and quotas

Supported AWS Regions

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

Guidelines

For best results:

- Use a lossless format, such as FLAC or WAV, with PCM 16-bit encoding.
- Use a sample rate of 16,000 Hz or higher.

Amazon Transcribe Medical may store your content to continuously improve the quality of its analysis models. See the [Amazon Transcribe Medical FAQ](#) to learn more. To request that we delete content that may have been stored by Amazon Transcribe Medical, open a case with AWS Support.

Quotas

You can request a quota increase for the following resources.

Resource	Default
Number of concurrent batch transcription jobs	250
Number of concurrent HTTP/2 or WebSocket requests	25
Total number of medical vocabularies per AWS account	100
Number of pending medical vocabularies	10

The following API limits can also be increased upon request:

API	Maximum transactions per second
StartMedicalTranscriptionJob	25
StartMedicalStreamTranscription	25
GetMedicalTranscriptionJob	30
DeleteMedicalTranscriptionJob	5
ListMedicalTranscriptionJobs	5
CreateMedicalVocabulary	10
UpdateMedicalVocabulary	10
DeleteMedicalVocabulary	5
GetMedicalVocabulary	20
ListMedicalVocabularies	5

Note

For information about requesting a quota increase, see [AWS service quotas](#) in the *AWS General Reference*.

The following quotas **cannot** be increased:

Description	Quotas
Maximum audio file length	14,400 seconds
Maximum audio file size	2 GB
Number of days that job records are retained	90
Number of channels for channel identification	2

Description	Quotas
Minimum audio file duration, in milliseconds (ms)	500

Amazon Transcribe Medical and interface VPC endpoints (AWS PrivateLink)

You can establish a private connection between your VPC and Amazon Transcribe Medical by creating an *interface VPC endpoint*. Interface endpoints are powered by [AWS PrivateLink](#), a technology that enables you to privately access Amazon Transcribe Medical APIs without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with Amazon Transcribe Medical APIs. Traffic between your VPC and Amazon Transcribe Medical does not leave the Amazon network.

Each interface endpoint is represented by one or more [Elastic Network Interfaces](#) in your subnets.

For more information, see [Interface VPC endpoints \(AWS PrivateLink\)](#) in the *Amazon VPC User Guide*.

Considerations for Amazon Transcribe Medical VPC endpoints

Before you set up an interface VPC endpoint for Amazon Transcribe Medical, ensure that you review [Interface endpoint properties and limitations](#) in the *Amazon VPC User Guide*.

Amazon Transcribe Medical supports making calls to all of its API actions from your VPC.

Creating an interface VPC endpoint for Amazon Transcribe Medical

You can create a VPC endpoint for the Amazon Transcribe Medical service using either the AWS Management Console or the AWS CLI. For more information, see [Creating an interface endpoint](#) in the *Amazon VPC User Guide*.

For batch transcription in Amazon Transcribe Medical, create a VPC endpoint using the following service name:

- `com.amazonaws.us-west-2.transcribe`

For streaming transcription in Amazon Transcribe Medical, create a VPC endpoint using the following service name:

- `com.amazonaws.us-west-2.transcribestreaming`

If you enable private DNS for the endpoint, you can make API requests to Amazon Transcribe Medical using its default DNS name for the AWS Region, for example, `transcribestreaming.us-east-2.amazonaws.com`.

For more information, see [Accessing a service through an interface endpoint](#) in the *Amazon VPC User Guide*.

Creating a VPC endpoint policy for Amazon Transcribe Medical streaming

You can attach an endpoint policy to your VPC endpoint that controls access to Amazon Transcribe Medical. The policy specifies the following information:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which actions can be performed.

For more information, see [Controlling access to services with VPC endpoints](#) in the *Amazon VPC User Guide*.

Example: VPC endpoint policy for Amazon Transcribe Medical streaming transcription actions

The following is an example of an endpoint policy for streaming transcription in Amazon Transcribe Medical. When attached to an endpoint, this policy grants access to the listed Amazon Transcribe Medical actions for all principals on all resources.

```
{
  "Statement": [
    {
      "Principal": "*",
      "Effect": "Allow",
      "Action": [
        "transcribe:StartMedicalStreamTranscription",
      ],
      "Resource": "*"
    }
  ]
}
```

Example: VPC endpoint policy for Amazon Transcribe Medical batch transcription actions

The following is an example of an endpoint policy for batch transcription in Amazon Transcribe Medical. When attached to an endpoint, this policy grants access to the listed Amazon Transcribe Medical actions for all principals on all resources.

```
{
  "Statement": [
    {
      "Principal": "*",
      "Effect": "Allow",
      "Action": [
        "transcribe:StartMedicalTranscriptionJob"
      ],
      "Resource": "*"
    }
  ]
}
```

Document history for Amazon Transcribe and Amazon Transcribe Medical

- **Latest documentation update:** 31 May 2022

The following table describes important changes in each release of Amazon Transcribe. For notification about updates to this documentation, you can subscribe to an RSS feed.

update-history-change	update-history-description	update-history-date
Section Update	Batch language identification can now identify multiple languages per audio file.	May 31, 2022
Guide Update	The Amazon Transcribe API Reference is now a standalone guide.	April 1, 2022
New chapter	A new comparison table for Amazon Transcribe, Amazon Transcribe Medical, and Amazon Transcribe Call Analytics is included.	March 21, 2022
New chapter	A new SDK code examples chapter is included.	March 21, 2022
Feature update	Call Analytics now provides call summarization.	March 21, 2022
Chapter update	The introductory chapter now showcases Amazon Transcribe use cases.	March 21, 2022
Chapter update	The Getting started chapter has been updated to be method-specific.	March 21, 2022
Chapter update	The Streaming chapter has been updated and restructured.	March 21, 2022
Feature update	Language identification now supports custom vocabularies and custom vocabulary filters with streaming transcriptions.	March 11, 2022
New event	There is a new event type: Vocabulary events.	February 7, 2022
Section update	Updates have been made to the custom vocabularies section.	January 20, 2022

New feature	Language identification can now be used with streaming transcriptions.	November 23, 2021
New feature	Language identification can now be used with custom language models, custom vocabularies, vocabulary filtering, and content redaction.	October 29, 2021
New feature	Amazon Transcribe now supports custom language models with streaming transcriptions.	October 20, 2021
New feature	Amazon Transcribe can now generate subtitles for your video files.	September 16, 2021
New feature	Amazon Transcribe now supports PII redaction and identification for streaming.	September 14, 2021
New feature	Amazon Transcribe now supports AWS KMS encryption context for an added level of security for your AWS account resources.	September 10, 2021
New languages	Amazon Transcribe now supports Afrikaans, Danish, Mandarin Chinese (Traditional), Thai, New Zealand English, and South African English.	August 26, 2021
New feature	Amazon Transcribe now supports resource tagging.	August 24, 2021
New feature	Amazon Transcribe now supports Call Analytics for batch transcription jobs.	August 4, 2021
New feature	Amazon Transcribe Medical now supports personal health information (PHI) identification in batch transcription jobs.	May 14, 2021
New feature	Amazon Transcribe now supports using custom vocabularies with batch custom language models.	May 12, 2021
New feature	Amazon Transcribe now supports partial result stabilization for streaming transcription.	May 11, 2021

New feature	Amazon Transcribe now supports Australian English, British English, Hindi, and US Spanish for custom language models.	March 19, 2021
New feature	Amazon Transcribe Medical now supports personal health information identification.	January 29, 2021
New feature	Amazon Transcribe Medical now supports channel identification.	December 21, 2020
New feature	Amazon Transcribe Medical now supports the HTTP/2 streaming protocol.	November 24, 2020
New feature	Amazon Transcribe Medical now supports the neurology medical specialty for streaming audio transcription.	November 24, 2020
New feature	Amazon Transcribe Medical now supports the urology medical specialty for streaming audio transcription.	November 24, 2020
New feature	Amazon Transcribe Medical now supports the radiology medical specialty for streaming audio transcription.	November 24, 2020
New feature	Amazon Transcribe Medical now supports the cardiology medical specialty for streaming audio transcription.	November 24, 2020
New feature	Amazon Transcribe Medical now supports the oncology medical specialty for streaming audio transcription.	November 24, 2020
New feature	Amazon Transcribe now supports OGG/OPUS and FLAC codecs for streaming audio transcription.	November 24, 2020
New languages	Amazon Transcribe adds support for Italian and German for streaming audio transcription.	November 4, 2020
AWS Region expansion	Amazon Transcribe is now available in the Frankfurt (eu-central-1) and London (eu-west-2).	November 4, 2020
New feature	Amazon Transcribe adds support for interface VPC endpoints in batch transcription.	October 9, 2020

New feature	Amazon Transcribe adds support for channel identification in streaming.	September 17, 2020
New feature	Amazon Transcribe adds support for automatic language identification in batch transcription.	September 15, 2020
New feature	Amazon Transcribe adds support for speaker identification in streaming.	August 19, 2020
New feature	Amazon Transcribe adds support for custom language models.	August 5, 2020
New feature	Amazon Transcribe adds support for interface VPC endpoints in streaming.	June 26, 2020
New feature	Amazon Transcribe adds support for vocabulary filtering in streaming.	May 20, 2020
New feature	Amazon Transcribe Medical adds support for custom vocabularies in both batch processing and streaming.	April 29, 2020
New feature	Amazon Transcribe Medical adds support for batch processing of audio files.	April 1, 2020
New feature	Amazon Transcribe adds support for automatically redacting personally identifiable information.	February 26, 2020
New feature	Amazon Transcribe adds support for creating a custom vocabulary of words to filter from a transcription.	December 20, 2019
New feature	Amazon Transcribe adds support for queueing transcription jobs.	December 19, 2019
New languages	Amazon Transcribe adds support for Gulf Arabic, Hebrew, Japanese, Malay, Swiss German, Telugu, and Turkish.	November 21, 2019
AWS Region expansion	Amazon Transcribe is now available in the Asia Pacific (Tokyo) (ap-northeast-1).	November 21, 2019
New feature	Amazon Transcribe adds support for alternative transcriptions.	November 20, 2019

New languages	Amazon Transcribe adds support for Dutch, Farsi, Indonesian, Irish English, Portuguese, Scottish English, Tamil, and Welsh English.	November 12, 2019
New language	Amazon Transcribe now supports streaming transcription for Australian English (en-AU).	October 25, 2019
AWS Region expansion	Amazon Transcribe is now available in the China (Beijing) (cn-north-1) and China (Ningxia) (cn-northwest-1).	October 9, 2019
New feature	Amazon Transcribe enables you to provide your own KMS key to encrypt your transcription output files. For more information, see the OutputEncryptionKMSKeyId parameter of the StartStreamTranscription API.	September 24, 2019
New languages	Amazon Transcribe adds support for Chinese (Mandarin), Simplified, Mainland China and Russian.	August 23, 2019
New feature	Amazon Transcribe adds support for streaming audio transcription using the WebSocket protocol.	July 19, 2019
New feature	AWS CloudTrail now records events for the StartStreamTranscription API.	July 19, 2019
AWS Region expansion	Amazon Transcribe is now available in the US West (N. California) (us-west-1).	June 27, 2019
New language	Amazon Transcribe adds support for Modern Standard Arabic.	May 28, 2019
New feature	Amazon Transcribe now transcribes numeric words into numbers for US English. For example, "forty-two" is transcribed as "42".	May 23, 2019
New language	Amazon Transcribe adds support for Hindi and Indian English.	May 15, 2019
New SDK	The AWS SDK for C++ now supports Amazon Transcribe.	May 8, 2019
New language	Amazon Transcribe adds support for Spanish.	April 19, 2019

AWS Region expansion	Amazon Transcribe is now available in the EU (Frankfurt) (eu-central-1) and Asia Pacific (Seoul) (ap-northeast-2).	April 18, 2019
New language	Amazon Transcribe adds support for streaming transcription in British English, French, and Canadian French.	April 5, 2019
New feature	The AWS SDK for Ruby V3 now supports Amazon Transcribe	March 25, 2019
New feature	Amazon Transcribe now enables you to create custom vocabularies, lists of specific words that you want Amazon Transcribe to recognize in your audio input.	March 25, 2019
New languages	Amazon Transcribe adds support for German and Korean.	March 22, 2019
New language	Amazon Transcribe now supports streaming transcription for US Spanish (es-US).	February 7, 2019
AWS Region expansion	Amazon Transcribe is now available in the South America (Sao Paulo) (sa-east-1).	February 7, 2019
AWS Region expansion	Amazon Transcribe is now available in the Asia Pacific (Mumbai) (ap-south-1), Asia Pacific (Singapore) (ap-southeast-1), EU (London) (eu-west-2), and EU (Paris) (eu-west3).	January 24, 2019
New languages	Amazon Transcribe adds support for French, Italian, and Brazilian Portuguese.	December 20, 2018
New feature	Amazon Transcribe now supports transcription of audio streams.	November 19, 2018
New languages	Amazon Transcribe adds support for Australian English, British English, and Canadian French.	November 15, 2018
AWS Region expansion	Amazon Transcribe is now available in Canada (Central) (ca-central-1) and Asia Pacific (Sydney) (ap-southeast-2).	July 17, 2018
New feature	You can now specify your own location to store the output from a transcription job.	July 11, 2018

New feature	Added AWS CloudTrail and Amazon CloudWatch Events integration.	June 28, 2018
New feature	Amazon Transcribe adds support for custom vocabularies.	April 4, 2018
New guide (p. 333)	This is the first release of the <i>Amazon Transcribe Developer Guide</i> .	November 29, 2017

AWS glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS General Reference*.