
Live Streaming with Automated Multi-Language Subtitling Implementation Guide



Live Streaming with Automated Multi-Language Subtitling: Implementation Guide

Copyright © 2021 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

Home	1
Overview	2
Cost	2
Architecture	2
Considerations	5
Amazon Transcribe streaming limits	5
Encoding profile	5
Amazon VPC limits	5
Supported languages	5
Changing translated languages	5
Changing transcribed languages	6
Live Streaming with Automated Multi-Language Subtitling Update	6
Regional deployment	6
Template	7
Deployment	8
Launch the stack	8
Security	11
Amazon CloudFront	11
Resources	12
Collection of Operational Metrics	13
Troubleshooting	14
Source Code	15
Revisions	16
Contributors	17
Notices	18

Live Streaming with Automated Multi-Language Subtitling

AWS Implementation Guide

AWS Solutions Builder Team

December 2020

This implementation guide describes architectural considerations and provides configuration steps for deploying the Live Streaming with Automated Multi-Language Subtitling solution in the Amazon Web Services (AWS) Cloud. It includes links to an [AWS CloudFormation](#) template that launches and configures the AWS services required to deploy this solution using AWS best practices for security and availability.

The guide is intended for IT infrastructure architects, administrators, and DevOps professionals who have practical experience with video streaming and architecting on the AWS Cloud.

Overview

Amazon Web Services (AWS) offers Live Streaming with Automated Multi-Language Subtitling, a solution that automatically generates multi-language subtitles for live streaming video content in real time. The solution is used only during the live event.

Adding subtitles to your live video content can help improve reach and access, exposing your content to a much larger audience. However, it can be a challenging and time-consuming process to add subtitles to a live stream. The process can become more difficult when multi-language subtitles are required.

The solution uses AWS Elemental MediaLive, and AWS Elemental MediaPackage to encode and package your content for adaptive bitrate streaming across multiple screens, and [AWS Lambda](#), [Amazon Transcribe Streaming](#), and [Amazon Translate](#) to convert the audio to text and generate captions in multiple languages.

This solution uses Machine Learning (ML) services for transcription and translation, without the need for a human stenographer. This enables you to automate the subtitling process and reach a wider range of customers. However, the quality of transcription and translation services are affected by numerous factors and could result in suboptimal output. This solution might require additional customization or human supervision for broadcast-grade applications.

This solution allows for a single input language and up to five translated caption languages. Similar to a stenographer, the subtitles are slightly time-delayed from the audio. This solution is optimized for two second HTTP Live Streaming (HLS) segments on [AWS Elemental MediaLive](#), results are unknown with different segment sizes and could result in poor user experience. This implementation might not be suitable as a replacement for a human stenographer, especially for broadcast applications where users are familiar with human-generated subtitles.

Cost

You are responsible for the cost of the AWS services used while running this subtitling solution. As of the date of publication, the cost for running this solution in the US East (N. Virginia) Region is approximately **\$3.00 per hour** for the AWS Elemental MediaLive and AWS Elemental MediaPackage. It is approximately **\$1.44 per hour per stream** for the English Amazon Transcribe Streaming subtitles. Each additional language is approximately **\$0.50 per hour per stream**. Using this solution with the default settings will cost approximately \$5.00 an hour. The size of your viewership is the largest cost factor. To estimate your costs depending on your viewership, review the [Amazon CloudFront pricing page](#). Prices are subject to change. For full details, refer the pricing webpage for each AWS service you will be using in this solution.

Architecture overview

Deploying this solution with the default parameters builds the following environment in the AWS Cloud.

Live Streaming with Automated Multi-Language Subtitling Implementation Guide
Architecture

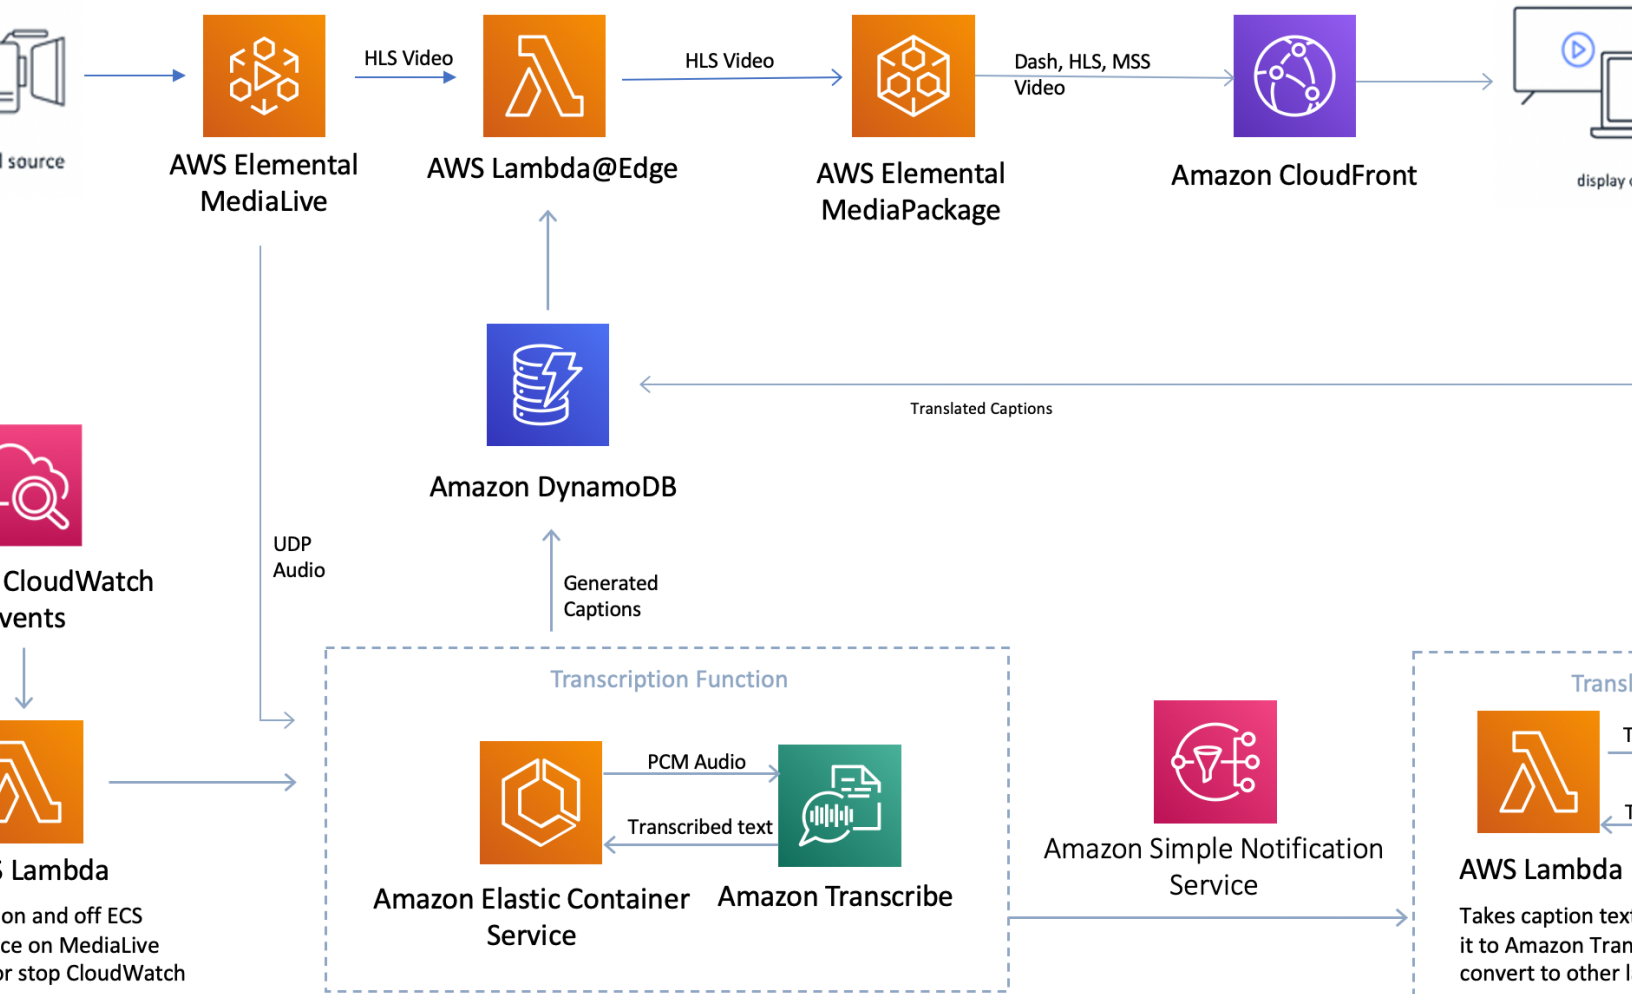


Figure 1: Live Streaming with Automated Multi-Language Subtitling solution architecture

The solution's AWS CloudFormation template deploys Live Streaming on AWS with additional services:

- [AWS Elemental MediaLive](#)
- [AWS Elemental MediaPackage](#)
- [Amazon CloudFront](#)
- [Amazon CloudWatch Events](#)
- [Amazon Simple Notification Service](#) (Amazon SNS)
- [Amazon Simple Storage Service](#) (Amazon S3)
- [Amazon Transcribe](#)
- [Amazon Elastic Container Service](#) (Amazon ECS)
- [Amazon Translate](#)
- [AWS Lambda](#)

The `TranslateFunction` Lambda function converts the transcribed text into different languages. Another AWS Lambda function is deployed to disable the transcription function when the AWS Elemental MediaLive channel stops, and enable it when the AWS Elemental MediaLive channel starts.

In order to get WebVtt subtitles into AWS Elemental MediaPackage, we use the `Lambda@Edge` function in Amazon CloudFront to insert subtitles into the HLS stream, which is sent from AWS Elemental MediaLive to AWS Elemental MediaPackage. Starting at the left of the diagram, Amazon MediaLive outputs HLS to an Amazon CloudFront endpoint. This endpoint passes through the video files, manifests, and only invokes the `Lambda@Edge` function for the WebVTT subtitle files passing from AWS Elemental MediaLive to AWS Elemental MediaPackage. Subtitles are inserted into these WebVTT files before they are passed on to Amazon MediaPackage.

AWS Elemental MediaLive outputs an audio-only User Datagram Protocol (UDP) stream to an Amazon ECS container. This container transmits the audio to Amazon Transcribe Streaming, which receives the text contained in the stream as asynchronous responses and writes each text response to an Amazon Dynamo DB table. This Amazon ECS container also sends Amazon SNS notifications to an Amazon Translate Lambda function, which creates translated subtitles that are written to the same Amazon Dynamo DB table.

Each WebVTT file invokes the `Lambda@Edge` function, which inserts the subtitles and then transmits them onto the MediaPackage. MediaLive provides the authentication headers.

MediaPackage ingests the files and packages them into formats that are delivered to four MediaPackage custom endpoints.

An Amazon CloudFront distribution is configured to use the MediaPackage custom endpoints as its origin. This CloudFront URL is what is provided to the viewers of the live stream.

Considerations

Amazon Transcribe streaming limits

Amazon Transcribe Streaming is used within the Amazon ECS container. The Amazon Transcribe Streaming quota is five concurrent streams and we recommend requesting a service limit increase for the number of Amazon Transcribe Streams. For more information on limits, refer to [Amazon Transcribe Limits](#). To request a limits increase, use the [Amazon Transcribe service limits increase form](#).

Encoding profile

This solution leverages the AWS Elemental MediaLive encoding profile from the Live Streaming on AWS solution. The Live Streaming on AWS solution includes the following encoding profile.

- **1080p profile:** 1080p@6000kbps, 720p@3000kbps, 480p@1500kbps, 240p@750kbps

Amazon VPC limits

This solution deploys a new Amazon Virtual Private Cloud (VPC) for the Amazon Transcribe ECS instance. If you plan to deploy more than once instance of this solution in one AWS Region, you may need to increase the Amazon VPC quota for your target Region. The default Amazon VPC limit is five per Region.

Supported languages

This solution currently supports English as the input audio language. For a list of supported translated output languages, refer to [Supported Languages and Language Codes](#) in the *Amazon Translate Developer Guide*.

Changing translated languages

To change the output languages, you must 1. update the caption output and 2. update the **Name Modifier**. If you add additional caption outputs to the MediaLive channel, you must add the language code to the SNSTriggerAWSTranslateLambda function as well. To change the translated language:

1. Log in to the [AWS MediaLive console](#).
2. Locate the appropriate channel and under **Channel Actions**, select **Edit Channel**. If the channel is already running, choose **Stop Channel** first.
3. Under **Output groups**, choose **Live (HLS)** and choose **Add output** to add additional translated output languages. Update the **Name Modifier** with the language code from the [Supported language codes](#) in the *Amazon Translate Developer Guide*. For example: The **Name Modifier** for Spanish is `_es`.
4. Choose **Update Channel**.
5. Navigate to the [AWS Lambda console](#).

6. Locate the `SNSTriggerAWSTranslateLambda` Lambda function, and update the `CAPTION_LANGUAGES` variable using the appropriate language code from the [Supported language codes](#) in the *Amazon Translate Developer Guide*. Use a comma to separate multiple languages. For example: `CAPTION_LANGUAGES: en, es, fr, de`.

Changing transcribed languages

This solution uses AWS Transcribe Streaming to transcribe the source language stored in AWS Elemental MediaLive. The default source language for transcriptions is English (en-US). If your source content is in a different language, change the `TranslateLanguage` input when launching your CloudFormation template. In addition, you must edit the AWS Elemental MediaLive channel **Output 5: english**. Modify the **Name Modifier** language code from `_en` to your selected language. For example, if your selected language is Spanish (es-US), update the Name Modifier with `_es` when deploying the CloudFormation stack. You can also change the language code and language description as well. For more information, refer to [Changing translated languages \(p. 5\)](#) section.

This solution currently only supports the 16 kHz Amazon Transcribe Streaming languages. For more information about supported transcription languages, refer to [Streaming Transcription](#) in the *Amazon Translate Developer Guide*.

Live Streaming with Automated Multi-Language Subtitling Update

If you have previously deployed this solution you cannot upgrade from a previous stack, you must redeploy the CloudFormation stack with the latest version of the solution's framework.

Regional deployment

This solution uses Amazon Translate, AWS Elemental MediaLive, AWS Elemental MediaPackage, Amazon Transcribe Streaming, and AWS Elemental MediaConnect which are currently available in specific AWS Regions only. Therefore, you must launch this solution in an AWS Region where these services are available. For the most current service availability by region, refer to [AWS Service offerings by Region](#).

AWS CloudFormation template

This solution uses AWS CloudFormation to automate the deployment of Live Streaming with Automated Multi-Language Subtitling in the AWS Cloud. It includes the following AWS CloudFormation template, which you can download before deployment:

[View
Template](#)

live-streaming-with-automated-multi-language-subtitling.template: Use this template to launch the solution and all associated components. The default configuration deploys the Live Streaming on AWS solution, AWS Lambda functions, Amazon Simple Storage Service (Amazon S3) buckets, Amazon Transcribe, and Amazon Translate. You can also customize the template based on your specific needs. For more information about the Live Streaming on AWS solution, refer to the [Architecture Overview](#) section in the *Live Streaming on AWS Implementation Guide*.

Automated deployment

Before you launch the automated deployment, please review the considerations discussed in this guide. Follow the step-by-step instructions in this section to configure and deploy the solution into your account.

Time to deploy: Approximately 20 minutes

Launch the stack

This automated AWS CloudFormation template deploys the Live Streaming with Automated Multi-Language Subtitling solution on the AWS Cloud.

Note

You are responsible for the cost of the AWS services used while running this solution. See the [Cost \(p. 2\)](#) section for more details. For full details, see the pricing webpage for each AWS service you will be using in this solution.

1. Log in to the AWS Management Console and click the button below to launch the `live-streaming-with-automated-multi-language-subtitling` AWS CloudFormation template.



You can also [download the template](#) as a starting point for your own implementation.

2. The template is launched in the US East (N. Virginia) Region by default. To launch this solution in a different AWS Region, use the region selector in the console navigation bar.

Note

This solution uses the Amazon Translate, Amazon Transcribe, AWS Elemental MediaLive, MediaPackage, and MediaConnect services, which are currently available in specific AWS Regions only. Therefore, you must launch this solution in an AWS Region where these services are available. For the most current service availability by region, see the [AWS service offerings by region](#).

3. On the **Select Template** page, verify that you selected the correct template and choose **Next**.
4. On the **Specify Details** page, assign a name to your solution stack.
5. Under **Parameters**, review the parameters for the template, and modify them as necessary.

This solution uses the following default values.

Parameter	Default	Description
Live Streaming Source		
Source Input Type	URL Pull	Specify the input type for AWS Elemental MediaLive: RTP_PUSH, RTMP_PUSH, RTMP_PULL, URL_PULL, or MEDIACONNECT
URL_PULL CONFIGURATION		

Live Streaming with Automated Multi-
Language Subtitling Implementation Guide
Launch the stack

Parameter	Default	Description
Source URL (REQUIRED)	https://d15an60oaeed9r.cloudfront.net/live_stream_v2/sports_reel_with_markers.m3u8	The primary source URL for the live feed. By default, this parameter contains the primary demo source URL.
Source Username	<Optional input>	If authentication is required to access the source, enter the username.
Source Password	<Optional input>	If authentication is required to access the source, enter the password.
RTP_PUSH / RTMP_PUSH CONFIGURATION		
Input Security Group CIDR Block	<Requires input>	For RTP and RTMP PUSH input types ONLY, specify the CIDR Block for the MediaLive SecurityGroup. Input security group restricts access to the input and prevents unauthorized third parties from pushing content into a channel that is associated with that input.
MediaConnect Source Type		
MediaConnect Flow ARN	<Optional input>	The primary source MediaConnect flow for the live feed. You can create the flow in the MediaConnect console. To provide redundancy, create the primary and secondary flows in different Availability Zones.
Subtitles Config Using Amazon Transcribe		
Transcribe Language	en-US, es-US, de-DE, ja-JP, pt-BR, it-IT	Specify the input language for Amazon Transcribe Streaming to use. This is for transcribing the audio of the video stream you send to AWS MediaLive. For more information, refer to the Changing transcribed languages (p. 6) .
Custom Vocabulary	<Optional input>	Specify a name of an existing custom vocabulary for Amazon Transcribe to use. For more information, refer to Custom Vocabularies in the <i>Amazon Transcribe Developer Guide</i> .

Parameter	Default	Description
Vocabulary Filter Name	<Optional input>	Specify a name of an existing vocabulary filter for Amazon Transcribe to use. For more information, refer to Custom Vocabulary Filters in the <i>Amazon Transcribe Developer Guide</i> .
Encoding Options		
Source Codec	AVC	Specify the codec for the source stream for MediaLive: AVC, HEVC, or MPEG2
Encoding Profile	1080	Specify the encoding profile to use with MediaLive.
Start MediaLive Channel	True	Choose whether to start the MediaLive channel when the solution is created.

6. Choose **Next**.
7. On the **Options** page, choose **Next**.
8. On the **Review** page, review and confirm the settings. Check the box acknowledging that the template creates AWS Identity and Access Management (IAM) resources.
9. Choose **Create stack** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation console in the Status column. You should receive a **CREATE_COMPLETE** status in approximately 20 minutes.

Note

In addition to the primary AWS Lambda function, this solution includes the `solution-helper` Lambda function, which runs only during initial configuration or when resources are updated or deleted.

When running this solution, you will see both Lambda functions in the AWS Lambda console, but only the primary Lambda function is regularly active. However, do not delete the `solution-helper` function because it is necessary to manage associated resources.

When you finish streaming, you can delete the solution's stack to ensure that you only pay for the infrastructure that you use. The `Lambda@Edge` function must be manually deleted when you delete the stack. To do this navigate to the AWS Lambda console in the Virginia (us-east-1) Region. The AWS Lambda has a name containing **EdgeCaptionInserter**. For more information, refer to [Lambda delete replicas](#).

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This shared model can reduce your operational burden as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. For more information about security on AWS, visit the [AWS Security Center](#).

Amazon CloudFront

This solution deploys a static website [hosted](#) in an Amazon S3 bucket. To help reduce latency and improve security, this solution includes an Amazon CloudFront distribution with an origin access identity, which is a special CloudFront user that helps restrict access to the solution's website bucket contents. For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#).

Additional Resources

AWS services

- [Amazon CloudFront](#)
- [AWS CloudFormation](#)
- [AWS Elemental MediaConnect](#)
- [AWS Elemental MediaLive](#)
- [AWS Elemental MediaPackage](#)
- [AWS Lambda](#)
- [Amazon Simple Storage Service](#)
- [Amazon Transcribe](#)
- [Amazon Translate](#)

AWS blogs

- [Connecting AWS Elemental Live On-Premises to AWS Media Services in the Cloud](#)
- [Connecting OBS Studio to AWS Media Services in the Cloud](#)
- [Connecting FFmpeg Using RTP to AWS Media Services in the Cloud](#)
- [Connecting FFmpeg Using RTMP to AWS Media Services in the Cloud](#)
- [Connecting VLC Media Playing Using RTP to AWS Media Services in the Cloud](#)

Tutorials and Workshops

- [AWS Live Streaming and Live-to-VOD Workshop](#)
- [Build a Secure Packager and Encoder Key Exchange \(SPEKE\) Reference Server](#)

Collection of operational metrics

This solution includes an option to send operational metrics to AWS. We use this data to better understand how customers use this solution and related products and services. When enabled, the following information is collected and sent to AWS when the AWS CloudFormation template is launched:

- **Solution ID:** The AWS solution identifier
- **Unique ID (UUID):** Randomly generated, unique identifier for each live streaming solution deployment
- **Timestamp:** Data-collection timestamp
- **Launch Data:** The AWS Region where the stack was launched
- **Source Input Type:** The input type you selected at launch
- **Source Input Codec:** The codec you selected at launch
- **Source Input Resolution:** The resolution you selected at launch
- **Source Code:** The code base you selected at launch

Note that AWS will own the data gathered via this survey. Data collection will be subject to the [AWS Privacy Policy](#). To opt out of this feature, modify the AWS CloudFormation template mapping section as follows:

```
AnonymousData:  
  SendAnonymousData:  
    Data: Yes
```

to

```
AnonymousData:  
  SendAnonymousData:  
    Data: No
```


Troubleshooting

If you are not getting subtitles when watching your live stream, follow these troubleshooting instructions:

1. Log in to the AWS Management console.
2. Navigate to the Amazon ECS console and locate the ECS container that contains the Stack Name you used to create your ECS service.
3. On the **Tasks** tab in the ECS service, choose the task name.
4. Choose the **log** tab. The logs listed can show exceptions, such as **LimitExceeded**. In this case you have to increase your Amazon Transcribe Streaming limit on your AWS account for the region you are running in.
5. If you are still not seeing subtitles, restart the ECS task by choosing **stop**. When you stop an ECS task a new task should start in a few minutes.

Source Code

Visit our [GitHub repository](#) to download the source files for this solution and to share your customizations with others.

Revisions

Date	Change	Location
March 2019	Initial release	
November 2020	Updated implementation guide to remove references to NodeJS.	
December 2020	Updated design considerations and cost estimate; for additional information about changes to v2.0, refer to the CHANGELOG.md file in the GitHub repository	

Contributors

The following individuals contributed to this document:

- Eddie Goynes
- Tom Gilman
- Leah Siddall

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS current product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied. AWS responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Live Streaming with Automated Multi-Language Subtitling is licensed under the terms of the of the Apache License Version 2.0 available at [Classless Inter-Domain Routing \(CIDR\)](#).