

Implementation Guide

Content Localization on AWS



Content Localization on AWS: Implementation Guide

Copyright © 2024 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

Solution overview	1
Why use customizations and human-in-the-loop?	2
Features and benefits	3
Use cases	4
Concepts and definitions	4
Architecture overview	7
AWS Well-Architected design considerations	8
Operational excellence	8
Security	9
Reliability	9
Performance efficiency	9
Cost optimization	9
Sustainability	10
Architecture details	11
AWS services in this solution	11
Media Insights on AWS solution	13
Media Insights on AWS content localization workflow	13
Operators in the default workflow	14
Web application	16
Web application authentication	17
OpenSearch Service Media Insights pipeline consumer	18
Plan your deployment	20
Supported AWS Regions	20
Cost	21
Sample cost table	21
Security	23
Server-side encryption	23
Amazon CloudFront	23
Quotas	24
Deploy the solution	25
Deployment process overview	25
CloudFormation template	26
Launch the stack	26
Monitor the solution with Service Catalog AppRegistry	29

Activate CloudWatch Application Insights	29
Confirm cost tags associated with the solution	31
Activate cost allocation tags associated with the solution	31
AWS Cost Explorer	32
Update the solution	33
Troubleshooting	34
Problem: Amazon CloudWatch Events bus permissions error	34
Resolution	34
Contact Support	34
Create case	34
How can we help?	35
Additional information	35
Help us resolve your case faster	35
Solve now or contact us	35
Uninstall the solution	36
Using the AWS Management Console	36
Using AWS Command Line Interface	36
Deleting the Amazon S3 buckets	36
Using the application	38
Access the web application	38
Identify the URL	38
Upload a video and run a workflow	38
Supported input types	38
Run a workflow	38
Work with the media collection	40
Browse the collection of processed media assets	41
Search the media collection	41
Use full-text queries	42
Search high confidence data	43
Search data from individual operators	43
Search related concepts across multiple operators	43
View source-language subtitles	44
Edit source language subtitles	45
Use corrections to source language subtitles to create an Amazon Transcribe custom vocabulary	45
View target-language subtitles	46

Download target-language WebVTT and SRT subtitles	47
Download target-language audio file	47
Edit target language subtitles	48
Use corrections to target language subtitles to create an Amazon Translate terminology	48
Create user for the application	49
Developer guide	52
Source code	52
Reference	53
Anonymized data collection	53
Option 1: Deploy Media Insights on AWS separately with metric reporting disabled	54
Option 2: Modify the source template used to deploy Media Insights on AWS as a nested stack	54
Related resources	55
Contributors	55
Revisions	57
Notices	60

Create subtitles for your video-on-demand content to reach more audiences, globally

Publication date: *November 2021* ([last update: December 2024](#))

Content Localization on AWS helps extend the reach of your video-on-demand (VOD) content by efficiently creating accurate multi-language subtitles using Amazon Web Services (AWS) Artificial Intelligence (AI) services. You can make manual corrections to the automatically created subtitles and use advanced AWS AI service customization features to improve the results of the content automation. This solution is built on [Media Insights on AWS](#), a framework that helps accelerate the development of serverless applications that process video, images, audio, and text with artificial intelligence services and multimedia services on AWS.

Localization is the process of taking video content that was created for audiences in one geography and transforming it to make it relevant and accessible to audiences in a new geography. Creating alternative language subtitle tracks is central to the localization process. With this solution, content owners can use a guided experience for automatically generating and correcting subtitles for videos in multiple languages using AWS AI services. Corrections made by solution users can be used as feedback to customize the results of AWS AI services for future workflows. This type of AI/ML (Machine Learning) workflow, which incorporates user corrections is referred to as *human-in-the-loop*.

Content localization workflows can make use of advanced customization features provided by Amazon Transcribe and Amazon Translate:

- [Amazon Transcribe custom vocabulary](#) - you provide Amazon Transcribe with a list of terms that are specific to your content and how you want the terms to be displayed in transcripts.
- [Amazon Transcribe custom language models](#) - you provide Amazon Transcribe with text data (training data) to improve transcription accuracy for industry-specific terms, acronyms, and other phrases that it might not otherwise recognize. This becomes a [text corpus](#) list of terms that are specific to your content and how you want the terms to be displayed in transcripts.
- [Amazon Translate custom terminologies](#) - you provide Amazon Translate with a list of terms or phrases in the source language content and specify how you want them to appear in the translated result.
- [Amazon Translate parallel data for active custom translation](#) - you provide Amazon Translate with a list of parallel phrases: the source language and the phrase translated the way you

want it. The *parallel data* customizes Amazon Translate models so they create more contextual translations based on the sample you provide.

Application users can manually correct the results of the automation at different points in the automated workflow and then invoke a new workflow to include their corrections in downstream processing. Corrections are tracked and can be used to update Amazon Transcribe custom vocabularies and Amazon Translate custom terminologies to improve future results.

Why use customizations and human-in-the-loop?

Automating the creation of translated subtitles using AI/ML helps to speed up the process of localization for your content, but there are still challenges to achieve the level of accuracy that is required for specific use cases. With natural language processing, many aspects of the content itself may determine the level of accuracy AI/ML analysis is capable of achieving. Some content characteristics that can impact transcription and translation accuracy include: domain specific language, speaker accents and dialects, new words recently introduced to common language, the need for contextual interpretation of ambiguous phrases, and correct translation of proper names. AWS AI services provide a variety of features to help customize the results of the machine learning to specific content. Therefore, the workflow in this application seeks to provide users with a guided experience to use these customization features as an extension of their normal editing workflow.

This implementation guide discusses architectural considerations and configuration steps for deploying the Content Localization in AWS solution in the 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 and developers who have practical experience architecting in the AWS Cloud.

Use this navigation table to quickly find answers to these questions:

If you want to . . .	Read . . .
Know the cost for running this solution The one-time estimated cost for processing a 30-minute video with five languages using the default settings in the US East (N. Virginia) Region is USD \$5.64 per month.	Cost

If you want to . . .	Read . . .
Understand the security considerations for this solution.	Security
Know how to plan for quotas for this solution.	Quotas
Know which AWS Regions are supported for this solution.	Supported AWS Regions
View or download the AWS CloudFormation template included in this solution to automatically deploy the infrastructure resources (the “stack”) for this solution.	AWS CloudFormation template
Access the source code and optionally use the AWS Cloud Development Kit (AWS CDK) to deploy the solution.	GitHub repository

Features and benefits

The Content Localization on AWS solution provides the following features:

Use a simple web interface

Upload and analyze videos, and work with automatically generated subtitles using a simple web-based user interface.

Leverage AWS AI services

Automatically extract valuable metadata from video files using [Amazon Rekognition](#), [Amazon Transcribe](#), [Amazon Translate](#), and [Amazon Comprehend](#).

Automate manual processes, but control the final result

You can review subtitles and make corrections within the application. Once you are satisfied with the subtitles, rerun the workflow using the corrected input to regenerate downstream results.

Customize AWS AI services to produce more accurate results

Use the application to generate Amazon Transcribe custom vocabularies and Amazon Translate custom terminologies using the corrections you make to the subtitles. Provide these customizations when you upload a video and configure the automated workflow.

Leverage Media Insights on AWS

Media Insights on AWS is a framework that makes it easier for developers to build serverless applications that process video, images, audio, and text with AI and multimedia services on AWS.

Integration with AWS Service Catalog AppRegistry and AWS Systems Manager Application Manager, a capability of AWS Systems Manager

This solution includes a Service Catalog AppRegistry resource to register the solution's CloudFormation template and its underlying resources as an application in both AWS [AWS Service Catalog AppRegistry](#) and [AWS Systems Manager Application Manager](#). With this integration, you can centrally manage the solution's resources.

Use cases

Content enrichment

Media and entertainment companies have the opportunity to reduce the cost of media storage while improving the search functionality and monetization of media asset libraries by migrating them to the cloud. [Content enrichment solutions on AWS](#) let you set up workflows to migrate media assets to AWS, offering options to analyze video, images, and audio to generate valuable metadata upon migration. This allows you to streamline the media migration process and improve the search, discovery, and monetization of content.

Media Intelligence

Expanding quantities of media-based information such as videos, images, and audio are difficult to codify, classify, and organize. Applying AI/ML to analyze and collect information from within media results in easier storage, recall, and movement of otherwise hard to organize media files. Media Intelligence allows users to more easily label, recall, and combine their media assets.

Concepts and definitions

This section describes key concepts and defines terminology specific to this solution:

Application

A logical group of AWS resources that you want to operate as a unit.

Workflow API

A REST interface to create, update, delete, and run workflows and operators, and monitor workflows.

Control plane

Includes the workflow API and state machines for workflows. Workflow state machines are composed of operators from the MI operator library. When operators within the state machine are run, they interact with the MI data plane to store and retrieve derived asset and metadata generated from the workflow.

Operators

Generated state machines that call AWS Lambda functions to perform media analysis or media transformation tasks.

Workflows

Generated state machines that run operators in sequence.

Data plane

Stores media assets and their associated metadata that are generated by workflows. Implement a consumer of the Kinesis data stream in the data plane to extract, transform, and load (ETL) data from the master MI data store to downstream databases that support the data access patterns required by end-user applications.

Data plane API

A REST interface to create, update, delete, and retrieve media assets and their associated metadata.

Data plane pipeline

Stores metadata for an asset that can be retrieved as a single block or pages of data using an object's `AssetId` and `Metadata` type. Writing data to the pipeline initiates a copy of the data to be stored in Kinesis Data Streams. This data stream is the interface that end-user applications can connect to use data stored in the MI data plane.

Data pipeline consumer


Changes to the data plane DynamoDB table are reflected in a Kinesis data stream. For each record in that data stream, data pipeline consumers perform the necessary extract, transform, and load

tasks needed to replicate data, such as media metadata, to the data stores used by external applications. These ETL tasks are entirely use-case dependent and therefore must be user-defined.

For a general reference of AWS terms, see the [AWS glossary](#).

Architecture overview

This section provides a reference implementation architecture diagram for the components deployed with this solution.



Content Localization on AWS architecture

This architecture depends on the Media Insights on AWS development framework, which must be deployed in the AWS account in order to deploy the solution. Media Insights on AWS can be deployed separately or together with this solution as an option. For a full description of the Media Insights on AWS architecture, refer to the [implementation guide](#).

The AWS CloudFormation template deploys the following infrastructure:

1. An instance of the Media Insights on AWS solution.

Note

If you already have a deployed instance of Media Insights on AWS that you would like to use, you can deploy this solution with an alternate template.

2. An [Amazon CloudFront](#) distribution to serve the solution's web application.

3. An [Amazon Simple Storage Service](#) (Amazon S3) web source bucket for hosting the static web application.
4. An [Amazon Cognito](#) user pool to provide a user directory.
5. An Amazon Cognito identity pool to provide federation with [AWS Identity and Access Management](#) (IAM) for authentication and authorization to the web application.
6. [Amazon API Gateway](#) endpoints for the Media Insights on AWS workflow API, the Media Insights on AWS data plane API and the [Amazon OpenSearch Service](#) API endpoint.
7. An [AWS Step Functions](#) workflow created by Media Insights on AWS. The content localization workflow consists of [AWS Lambda](#) functions that run jobs in [Amazon Transcribe](#), [Amazon Translate](#), [AWS Elemental MediaConvert](#), and [Amazon Polly](#). These Lambda functions also interact with the Media Insights on AWS data plane API to store and retrieve media objects and metadata returned by media analysis jobs. The workflow can also optionally run [Amazon Rekognition](#) and [Amazon Comprehend](#) to provide additional analysis of the input.
8. A Lambda function to extract, transform, and load media metadata from the MI data pipeline into an Amazon OpenSearch Service cluster. This lambda function is invoked by the Media Insights on AWS data plane DynamoDB stream whenever asset metadata is modified in the Media Insights on AWS data plane.
9. An Amazon OpenSearch Service cluster to index media metadata.

AWS Well-Architected design considerations

We designed this solution with best practices from the [AWS Well-Architected Framework](#), which helps customers design and operate reliable, secure, efficient, and cost-effective workloads in the cloud.

This section describes how the design principles and best practices of the Well-Architected Framework were applied when building this solution.

Operational excellence

This section describes how we architected this solution using the principles and best practices of the [operational excellence pillar](#).

The Content Localization on AWS solution pushes metrics to [Amazon CloudWatch](#) at various stages to provide observability into the infrastructure, Lambda functions, AI services, Amazon Simple Storage Service (Amazon S3) buckets, and the rest of the solution components.

Security

This section describes how we architected this solution using the principles and best practices of the [security pillar](#).

AWS highly recommends that customers encrypt sensitive data in transit and at rest. This solution automatically encrypts media files and metadata at rest with Amazon S3 Server-Side Encryption (SSE). The Media Insights on AWS solution's Amazon Simple Notification Service (Amazon SNS) topics and Amazon DynamoDB tables are also encrypted at rest using SSE.

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 CloudFront user that provides public access to the solution's website bucket contents.

Reliability

This section describes how we architected this solution using the principles and best practices of the [reliability pillar](#).

Content Localization on AWS uses AWS serverless services wherever possible (for example, Lambda, API Gateway, Amazon S3, and DynamoDB) to ensure high availability and quick recovery from service failure.

Performance efficiency

This section describes how we architected this solution using the principles and best practices of the [performance efficiency pillar](#).

Content Localization on AWS uses serverless architecture and along with Media Insights on AWS, can be launched in any AWS Region that supports the AWS services used in the solution such as: AWS Lambda, Amazon API Gateway, Amazon S3, Amazon Rekognition, Amazon Translate, Amazon Transcribe, Amazon Comprehend, Amazon Polly, and AWS Elemental MediaConvert.

This solution is automatically tested and reviewed by solutions architects and subject matter experts for areas to experiment and improve.

Cost optimization

This section describes how we architected this solution using the principles and best practices of the [cost optimization pillar](#).

Content Localization on AWS uses serverless architecture; therefore, customers only get charged for what they use.

The solution is a modular framework that allows users to configure and tailor their own media workflows and use only the AWS services that they need.

Sustainability

This section describes how we architected this solution using the principles and best practices of the [sustainability pillar](#).

Content Localization on AWS uses managed and serverless services to minimize the environmental impact of the backend services. A critical component for sustainability provided by the solution is maximizing the usage of the AWS AI services. The serverless design of this solution (using Lambda, API Gateway, Amazon S3, and DynamoDB) are aimed at reducing carbon footprint compared to the footprint of continually operating on-premises servers.

Architecture details

The Content Localization on AWS solution consists of a number of components that are deployed to AWS using CloudFormation nested stacks. The components include the Media Insights on AWS solution, a Media Insights on AWS workflow, a web application, web application authentication components, and an Amazon OpenSearch Service MI custom pipeline consumer.

AWS services in this solution

AWS service	Description
Amazon API Gateway	Core. Entry point to interact with the control and data plane APIs where requests to create, read, update, delete, or run workflows begin, or data retrieval begin.
Amazon DynamoDB	Core. Stores workflow-related data, such as state machine definitions for operators, workflow configurations, and workflow run statuses.
AWS Identity and Access Management (IAM)	Core. Grants the solution's AWS Lambda function access to create Regional resources.
AWS Lambda	Core. Supports workflow orchestration, operators executions, and store workflow results.
Amazon Simple Notification Service	Core. Supports workflow execution rate limits.
Amazon Simple Queue Service	Core. Supports workflow execution rate limits.
Amazon Simple Storage Service	Core. Resource used for storing input user media and output transformed media by the workflow.

AWS service	Description
AWS Step Functions	Core. Creates a resource when a user defines a new workflow using the workflow API.
Amazon Comprehend	Supporting. Can be integrated into workflows to find key phrases in text and references to real-world objects, dates, and quantities in text.
Amazon Elemental MediaConvert	Supporting. Can be integrated into workflows to transcode input video into MPEG-4 format and generate thumbnails.
Amazon Kinesis Data Streams	Supporting. Used to stream data changes reflected in DynamoDB which consumers can ETL data from the data store to downstream services.
Amazon Polly	Supporting. Can be integrated into workflows to turn input text into speech.
Amazon Rekognition	Supporting. Can be integrated into workflows for Celebrity Recognition, Content Moderation, Face Detection, Face Search, Label Detection, Person Tracking, Shot, Text, and Technical Cue Detection.
AWS Systems Manager	Supporting. Provides application-level resource monitoring and visualization of resource operations and cost data.
Amazon Translate	Supporting. Can be integrated into workflows to translate input text.
Amazon Transcribe	Supporting. Can be integrated into workflows to create SRT or VTT caption files from video transcripts. It can also convert input audio to text.

AWS service	Description
AWS X-Ray	Supporting. Provides debugging tools for the Media Insights on AWS application.

Media Insights on AWS solution

Media Insights on AWS is deployed in one of two modes, depending on which option you choose:


1. As a nested stack from the `deployment/content-localization-on-aws.yaml` template.
2. As a standalone stack whose name is provided as a parameter input to the `deployment/content-localization-on-aws-use-existing-mie-stack.yaml` template.

Media Insights on AWS provides services for creating and running the content localization workflow, and for storage and retrieval of the media objects and metadata that are generated by the workflow for each input video (asset). At deployment time, this solution integrates with Media Insights on AWS to create a workflow using the MI operator library and MI workflow custom resources. At runtime, the Content Localization on AWS solution integrates with Media Insights on AWS through its REST APIs to run the content localization workflow, and store and retrieve media objects and analysis outputs that are created by the workflow. Finally, the solution integrates with Media Insights on AWS through DynamoDB stream events emitted by the MI data pipeline as workflows run to store analysis results in Amazon OpenSearch Service.

Media Insights on AWS content localization workflow

Template for nested stack: `deployment/content-localization-on-aws-video-workflow.yaml`

The `ContentLocalizationWorkflow` MI workflow orchestrates the analysis and application logic to automatically generate multi-language video subtitles. The workflow is composed of Media Insights Engine operators as shown in the following diagram:



Content localization workflow diagram

Like all Media Insights on AWS workflows, the content localization workflow is composed of a number of stages that are run sequentially by Media Insights on AWS when the application calls the `POST /workflow/execution` API. Each stage contains Media Insights on AWS operators that are run in parallel.

The workflow uses *MI workflow reprocessing* to update selected outputs of the workflow for an existing asset after application users edit and save source and target language subtitles. By only reprocessing dependent, downstream operators, the solution can save the cost of rerunning operators whose outputs will not change.

Note

The workflow entry points used by the Content Localization on AWS application are shown by the green dots in the diagram.

Operators in the default workflow

Media Insights on AWS operators that run within the content localization workflow analyze and/or transform the input video. Operators write analysis results (metadata about the input video) to the Media Insights on AWS data pipeline, which stores raw output in the Amazon S3 `DataplaneBucket` bucket and makes the data available to downstream Media Insights

on AWS data pipeline consumers. The solution employs one built-in pipeline consumer, an Amazon OpenSearch Service resource. The web application retrieves results of operators run in the workflow using the `MI DataplaneApiEndpoint` or the Amazon OpenSearch Service API. Transformed media outputs (such as [WebVTT](#) subtitle tracks) are also stored using the `MI Dataplane` API. Output media objects are stored in the `MI Amazon S3 Dataplane` bucket. The object paths can also be retrieved using the `MI Dataplane` API for the operator.

Both the solution's web application and the content localization workflow can retrieve the result of any operator that has already run by calling the `GET /metadata/<assetid>/<operator-name>` `MI DataplaneApiEndpoint`.

Pre-built Media Insights on AWS operators:

- `MediaInfo` - performs analysis on the video package and provides information about the format of the video, including the number and types of tracks, encoding formats for tracks, and other structural information.
- `Thumbnail` - uses AWS MediaConvert to generate thumbnail images, transcode the video into a uniformly formatted proxy format that will be used as input to downstream operators, and create an audio-only file from the input.
- `TranscribeVideo` - generates a transcript of the spoken audio in the video using Amazon Transcribe.
- `WebCaptions` - converts the transcript generated by Amazon Transcribe into subtitle blocks stored in JSON format that contains the original transcript word-level timeseries plus subtitles.

The `WebCaptions` JSON metadata for items is stored in the following format:

```
{
  "WebCaptions": [
    {
      "start": "0.04",
      "end": "0.69",
      "caption": "En Austin. ",
      "sourceCaption": "In Austin."
    },
    {
      "start": "0.7",
      "end": "4.66",
      "caption": "Es 60\u00b0 con rel\u00e9jalo. ",
      "sourceCaption": "It's 60\u00b0 with a chance of..."
    }
  ],
}
```

```
...  
  ]  
}
```

WebCaptions data structure attributes:

start – the start time of the caption relative to the start of the video.

end – the end time of the caption relative to the start of the video.

caption – the caption for this translation including the most recent edits made through the application.

sourceCaption – the original caption generated from the transcript of the source video.

- **TranslateWebcaptions** - Generate a collection of target-language subtitles in JSON format, from the source-language subtitles. Items in the collection contain the language code and the language-specific Operator Name that can be used to retrieve translated WebCaptions output for each language in the collection.
- **WebToVTTCaptions** - Generates a collection of target-language subtitles in WebVTT format from the source-language subtitles. Items in the collection contain the language codes and object paths of each available subtitle track.
- **WebToSRTCaptions** - Generates a collection of target-language subtitles in [SRT format](#) from the source-language subtitles. Items in the collection contain the language codes and object paths of each available subtitle tracks.
- **PollyWebcaptions** - generates a collection of stand-alone audio files for all languages included in the workflow. Audio timing is not synchronized to the video. It is meant to be used in listen-only mode (for example, in situations where there is not enough bandwidth to play the video). Contains the language codes and S3 object path of each available audio track.

Web application

Template for nested stack: `deployment/content-localization-on-aws-web.yaml`




Web application component

The Content Localization on AWS solution features a simple static web application hosted in Amazon S3 for uploading, analyzing, and browsing video collections and creating subtitles for them. The deployed application accesses MI-created Amazon API Gateway endpoints for running the Media Insights on AWS workflow API and the MI DataPlaneApiEndpoint for asset storage and retrieval. It accesses the Amazon OpenSearch Service API Gateway endpoint for search and metadata retrieval.

Web application authentication

Template for nested stack: `deployment/content-localization-on-aws-auth.yaml`




Web application authentication component

The Content Localization on AWS web application uses Amazon Cognito [user pools](#) and [identity pools](#) for user authentication. The web application uses [AWS Amplify](#) prebuilt UI components for interacting with the authentication services.

When authenticated users upload files through the application, the files are stored in private folders that correspond to their unique Amazon Cognito identifier to ensure fine-grained access control using AWS Identity and Access Management (IAM) policies.

Amazon OpenSearch Service Media Insights on AWS pipeline consumer

Template for nested stack: `deployment/content-localization-on-aws-opensearch.yaml`



Amazon OpenSearch Service Media Insights on AWS pipeline consumer

This solution supports full-featured search of the metadata generated by the ContentLocalizationWorkflow workflow using Amazon OpenSearch Service. It indexes the metadata generated from its Media Insights on AWS operators in OpenSearch Service by attaching an OpenSearch Service consumer to the Media Insights on AWS data pipeline. The Lambda OpenSearchConsumer Lambda function is invoked by the Media Insights on AWS data plane DynamoDB stream whenever metadata for a Media Insights on AWS operator is stored to the Media Insights on AWS data plane. The Lambda function accesses the new or updated metadata and stores it to an OpenSearch Service index in the OpenSearch Service ESDomain instance.

The Content Localization on AWS application uses [Apache Lucene](#) queries to Amazon OpenSearch Service to support direct search by end users on the **Collection** page and internally to provide fine-grained search for Amazon Rekognition operators on the **Analyze** page.

Plan your deployment

This section describes the [cost](#), [security](#), [Regions](#), and [quota](#) considerations prior to deploying the solution.

Supported AWS Regions

This solution uses the Amazon Translate, Amazon Comprehend, and Amazon Transcribe, Amazon Translate, and Amazon Polly services, which are currently available in specific AWS Regions only. In addition, specific features of these services may also be available in limited Regions. Therefore, you must launch this solution in an AWS Region where these services are available. For the most current availability by Region, refer to the [AWS Regional Services List](#).

Content Localization on AWS is supported in the following AWS Regions:

Region name	
US East (Ohio)	Asia Pacific (Sydney)
US East (N. Virginia)	Asia Pacific (Singapore)
US West (Northern California)*	Canada (Central)
US West (Oregon)	Europe (Frankfurt)
Asia Pacific (Mumbai)	Europe (Ireland)
Asia Pacific (Seoul)	Europe (London)
Asia Pacific (Tokyo)	Europe (Paris)*

Note

*One or both of the optional services (Amazon Comprehend and Amazon Rekognition) are not available in this Region.

Cost

You are responsible for the cost of the AWS services used while running this solution. The primary cost factors are from using Amazon OpenSearch Service, Amazon Transcribe, Amazon Translate, and Amazon Polly.

Video processing using the default transcription and translation workflow and translating to five languages incurs a one-time cost of about \$0.17 per minute of video duration, but this can vary between \$0.12 per minute and \$0.60 per minute depending on the video content, the number of translations you are generating, and whether customization features are turned on for Amazon Transcribe and Amazon Translate. The recurring costs for data storage, search, and other services amount to about \$3.00 per day. As of this revision, the one-time cost for processing a 30-minute video with five languages using the default settings is **\$5.64 per day**.

After a video is uploaded into the solution, the costs for processing are a one-time expense. However, data storage costs occur daily. Prices are subject to change. For full details, refer to the pricing web page for each AWS service used in this solution.

Sample cost table

The following table provides a sample cost breakdown for deploying this solution with the default parameters in the US East (N. Virginia) Region for one month.

You are responsible for the cost of the AWS services used while running this solution. The primary cost factors are from using Amazon OpenSearch Service (successor to Amazon Elasticsearch Service), Amazon Transcribe, Amazon Translate, and Amazon Polly.

The following examples are for a 30-minute video processed with workflows configured to create translations for five languages. For these examples, we'll use public pricing from US East (N. Virginia) Region.

Example 1: Default subtitles workflow only, no customizations for Amazon Transcribe and Translate

AWS service	Dimensions	Cost [USD]
Amazon Transcribe	30 minutes of standard audio transcription	\$0.72 / video

AWS service	Dimensions	Cost [USD]
Amazon Translate	30,500 characters, 5 languages	\$2.29 / video
Amazon Polly	30,500 characters, 5 languages	\$0.60 / video
AWS Elemental MediaConvert	30 minutes, basic tier (<= 30 fps) at SD rate	\$0.36 / video
Amazon OpenSearch Service	t3.small.search instance	\$1.39 / day
Amazon Kinesis Data Streams	Shard hours	\$0.28 / day
Total cost for processing one video, including daily costs:		\$5.64

Example 2: Custom vocabulary, custom terminology, active custom translation, including all optional computer vision operators

AWS service	Dimensions	Cost [USD]
Amazon Transcribe	30 minutes of standard audio transcription with 30 minutes custom language	\$0.90 / video
Amazon Translate	30,500 characters, 5 languages with Active Customer Translation Batch	\$9.15 / video
Amazon Polly	30,500 characters, 5 languages	\$0.60 / video
MediaConvert	30 minutes, basic tier (<= 30 fps) at SD rate	\$0.36 / video
Amazon Rekognition	30 minutes of video with face detection	\$3.00 / video

AWS service	Dimensions	Cost [USD]
Amazon OpenSearch Service	t3.small.search instance	\$1.97 / day
Amazon Kinesis Data Streams	Shard hours	\$0.56 / day
Total cost for processing one video, including daily costs:		\$16.54

We recommend creating a [budget](#) through [AWS Cost Explorer](#) to help manage costs. Prices are subject to change. For full details, see the pricing webpage for each AWS service used in this solution.

Security

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

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

Server-side encryption

AWS highly recommends that customers encrypt sensitive data in transit and at rest. This solution automatically encrypts media files and metadata at rest with Amazon Simple Storage Service (Amazon S3) Server-Side Encryption (SSE). The Media Insights Engine solution's Amazon Simple Notification Service (Amazon SNS) topics and Amazon DynamoDB tables are also encrypted at rest using SSE

Amazon CloudFront

This solution deploys a static website [hosted](#) in an Amazon Simple Storage Service (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 CloudFront user that provides public access to the solution's website bucket contents. For more information, refer to [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) in the *Amazon CloudFront Developer Guide*.

Quotas

As of the most recent revision, the Content Localization on AWS solution limits the size of files that can be uploaded using the web application to 3 GB.

Each run of the `ContentLocalizationWorkflow` workflow creates an Amazon Translate batch job for each target language chosen in the form on the **Upload** page. Therefore, the total number of jobs run concurrently by the application equals the sum of the number of languages selected by each concurrent job. When a workflow is started that exceeds the Amazon Translate batch jobs, the workflow displays a status of `ERROR` and the Step Functions run result for the `TranslateWebCaptions` operator contains a message that Amazon Translate encountered a limit exceeded error. For details, refer to the [Amazon Translate guidelines and quotas](#) page and capacity provisioned in your account.

Deploy the solution

This solution uses [AWS CloudFormation templates and stacks](#) to automate its deployment. The CloudFormation template specifies the AWS resources included in this solution and their properties. The CloudFormation stack provisions the resources that are described in the template.

Note

If you don't want to use the latest version of Media Insights on AWS or have a customized Media Insights on AWS instance that you would like to use with this solution, go to the GitHub and download the `deployment/content-localization-on-aws-use-existing-mie-stack.yaml` template. This automated deployment uses the `deployment/content-localization-on-aws.yaml` template, which deploys the latest version of Media Insights on AWS.

Deployment process overview

Follow the step-by-step instructions in this section to configure and deploy the solution into your account.

Time to deploy: Approximately 20 minutes

Before you launch the solution, review the [Cost](#), [Architecture](#), [Security](#), and other considerations discussed earlier in this guide.

Important

This solution includes an option to send anonymized operational metrics to AWS. We use this data to better understand how customers use this solution and related services and products. AWS owns the data gathered through this survey. Data collection is subject to the [AWS Privacy Notice](#).

To opt out of this feature, download the template, modify the AWS CloudFormation mapping section, and then use the AWS CloudFormation console to upload your updated template and deploy the solution. For more information, see the [Anonymized data collection](#) section of this guide.

AWS CloudFormation template

You can download the CloudFormation template for this solution before deploying it.

[View template](#)

localization-on-aws.template - Use this template to launch the solution and all associated components. The default configuration deploys the core and supporting services found in the [AWS services in this solution](#) section, but you can customize the template to meet your specific needs .

Note

The `content-localization-on-aws` template deploys all nested templates associated with the solution's components. Refer to the [Architecture details](#) section for the name of each nested template and the associated resources deployed.

This AWS CloudFormation template deploys Content Localization on AWS in the AWS Cloud. You must meet the following prerequisites before launching the stack:

Note

If you have previously deployed this solution, see [Update the solution](#) for update instructions.

Launch the stack

This automated AWS CloudFormation template deploys Content Localization on AWS.

Note

You are responsible for the cost of the AWS services used while running this solution. For more details, visit the [Cost](#) section in this guide, and refer to the pricing webpage for each AWS service used in this solution.

1. Sign in to the AWS Management Console and select the button to launch the content-localization-on-aws.template AWS CloudFormation template.

Launch solution

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

Note


This solution uses the Amazon Translate, Amazon Comprehend, and Amazon Transcribe, Amazon Translate and Amazon Polly services, which are currently available in specific AWS Regions only. In addition, specific features of these services may also be available in limited Regions. Therefore, you must launch this solution in an AWS Region where these services are available. For the most current availability by Region, refer to the [AWS Regional Services List](#).

3. On the **Create stack** page, verify that the correct template URL is in the **Amazon S3 URL** text box and choose **Next**.
4. On the **Specify stack details** page, assign a name to your solution stack. For information about naming character limitations, refer to [IAM and STS Limits](#) in the *AWS Identity and Access Management User Guide*.
5. Under **Parameters**, review the parameters for this solution template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
AdminEmail	<i><Requires input></i>	Email address for the MI administrator. This user receives an email with a temporary password to the web application once the AWS CloudFormation template has launched.
OpenSearchNodeSize	t3.small.search	The node type to be provisioned for the Amazon

Parameter	Default	Description
		OpenSearch Service. The default value for this parameter is <code>t3.small.search</code> . For information about which instance type is appropriate for your use case, refer to Sizing Amazon OpenSearch Service domains .

6. Choose **Next**.
7. On the **Configure stack options** page, choose **Next**.
8. On the **Review** page, review and confirm the settings. Check the box acknowledging that the template will create AWS Identity and Access Management (IAM) resources and may require an AWS CloudFormation capability.

 **Note**

This solution may require an AWS CloudFormation capability: `CAPABILITY_AUTO_EXPAND`, which is a parameter that supports the use of macros. For information about this AWS CloudFormation capability, refer to [CreateStack](#) in the *AWS CloudFormation API Reference*.

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.


Monitor the solution with Service Catalog AppRegistry

This solution includes a Service Catalog AppRegistry resource to register the CloudFormation template and underlying resources as an application in both [Service Catalog AppRegistry](#) and [AWS Systems Manager Application Manager](#).

AWS Systems Manager Application Manager gives you an application-level view into this solution and its resources so that you can:

- Monitor its resources, costs for the deployed resources across stacks and AWS accounts, and logs associated with this solution from a central location.
- View operations data for the resources of this solution (such as deployment status, CloudWatch alarms, resource configurations, and operational issues) in the context of an application.

The following figure depicts an example of the application view for the solution stack in Application Manager.



Solution stack in Application Manager

Activate CloudWatch Application Insights

1. Sign in to the [Systems Manager console](#).

2. In the navigation pane, choose **Application Manager**.
3. In **Applications**, search for the application name for this solution and select it.

The application name will have App Registry in the **Application Source** column, and will have a combination of the solution name, Region, account ID, or stack name.

4. In the **Components** tree, choose the application stack you want to activate.
5. In the **Monitoring** tab, in **Application Insights**, select **Auto-configure Application Insights**.

The screenshot shows the AWS Application Insights Monitoring interface. At the top, there is a navigation bar with tabs: Overview, Resources, Provisioning, Compliance, **Monitoring**, OpsItems, Logs, Runbooks, and Cost. Below the navigation bar, the main content area is titled "Application Insights (0) Info" and includes a toggle for "View Ignored Problems", an "Actions" dropdown, and an "Add an application" button. A search bar contains the text "Find problems". To the right of the search bar are filters for "Last 7 days", a refresh button, and pagination controls showing "1" of 1 items. Below the search bar is a table header with columns: Problem su..., Status, Severity, Source, Start time, and Insights. The main content area displays a message: "Advanced monitoring is not enabled. When you onboard your first application, a service-linked role (SLR) is created in your account. The SLR is predefined by CloudWatch Application Insights and includes the permissions the service requires to monitor AWS services on your behalf." At the bottom of this message is a button labeled "Auto-configure Application Insights".


Monitoring for your applications is now activated and the following status box appears:

The screenshot shows the AWS Application Insights Monitoring interface after successful activation. The navigation bar and top controls are the same as in the previous screenshot. The main content area now displays a green-bordered status box with a checkmark icon and the text: "Application monitoring has been successfully enabled. It will take some time to display any results. Please use the refresh button to view results." The rest of the interface, including the search bar, filters, and table header, remains the same.

Confirm cost tags associated with the solution

After you activate cost allocation tags associated with the solution, you must confirm the cost allocation tags to see the costs for this solution. To confirm cost allocation tags:

1. Sign in to the [Systems Manager console](#).
2. In the navigation pane, choose **Application Manager**.
3. In **Applications**, choose the application name for this solution and select it.
4. In the **Overview** tab, in **Cost**, select **Add user tag**.



5. On the **Add user tag** page, enter `confirm`, then select **Add user tag**.

The activation process can take up to 24 hours to complete and the tag data to appear.

Activate cost allocation tags associated with the solution

After you confirm the cost tags associated with this solution, you must activate the cost allocation tags to see the costs for this solution. The cost allocation tags can only be activated from the management account for the organization.

To activate cost allocation tags:

1. Sign in to the [AWS Billing and Cost Management and Cost Management console](#).
2. In the navigation pane, select **Cost Allocation Tags**.
3. On the **Cost allocation tags** page, filter for the AppManagerCFNStackKey tag, then select the tag from the results shown.
4. Choose **Activate**.

AWS Cost Explorer

You can see the overview of the costs associated with the application and application components within the Application Manager console through integration with AWS Cost Explorer. Cost Explorer helps you manage costs by providing a view of your AWS resource costs and usage over time.

1. Sign in to the [AWS Cost Management console](#).
2. In the navigation menu, select **Cost Explorer** to view the solution's costs and usage over time.

Update the solution

If you have previously deployed the solution, follow this procedure to update the Content Localization on AWS CloudFormation stack to get the latest version of the solution's framework.

1. Sign in to the [AWS CloudFormation console](#), select your existing Content Localization on AWS CloudFormation stack, and select **Update**.
2. Select **Replace current template**.
3. Under **Specify template**:
 - a. Select **Amazon S3 URL**.
 - b. Copy the link of the latest template.
 - c. Paste the link in the **Amazon S3 URL** box.
 - d. Verify that the correct template URL shows in the **Amazon S3 URL** text box, and choose **Next**. Choose **Next** again.
4. Under **Parameters**, review the parameters for the template and modify them as necessary. For details about the parameters, see [Launch the stack](#).
5. Choose **Next**.
6. On the **Configure stack options** page, choose **Next**.
7. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.
8. Choose **View change set** and verify the changes.
9. Choose **Update 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 UPDATE_COMPLETE status in approximately 10 minutes.

Troubleshooting

If these instructions don't address your issue, [Contact AWS Support](#) provides instructions for opening an Support case for this solution.

Problem: Amazon CloudWatch Events bus permissions error

If during spoke stack deployment, you received a **CREATE_FAILED** message for the TAWarnRule and/or the TSErrorRule, verify that the CloudWatch Event Bus in the primary account allows the spoke account to send events to the primary account.

Resolution

Update the primary stack with the secondary account ID or complete the following task:

1. In the primary account, navigate to the [CloudWatch console](#).
2. In the navigation pane, select **Event Buses**.
3. Select **Add Permissions**.
4. For **Principal**, enter the applicable secondary account ID.
5. Select the **Everybody(*)** checkbox.
6. Choose **Add**.

Contact Support

If you have [AWS Developer Support](#), [AWS Business Support](#), or [AWS Enterprise Support](#), you can use the Support Center to get expert assistance with this solution. The following sections provide instructions.

Create case

1. Sign in to [Support Center](#).
2. Choose **Create case**.

How can we help?

1. Choose **Technical**.
2. For **Service**, select **Solutions**.
3. For **Category**, select **Other Solutions**.
4. For **Severity**, select the option that best matches your use case.
5. When you enter the **Service**, **Category**, and **Severity**, the interface populates links to common troubleshooting questions. If you can't resolve your question with these links, choose **Next step: Additional information**.

Additional information

1. For **Subject**, enter text summarizing your question or issue.
2. For **Description**, describe the issue in detail.
3. Choose **Attach files**.
4. Attach the information that AWS Support needs to process the request.

Help us resolve your case faster

1. Enter the requested information.
2. Choose **Next step: Solve now or contact us**.

Solve now or contact us

1. Review the **Solve now** solutions.
2. If you can't resolve your issue with these solutions, choose **Contact us**, enter the requested information, and choose **Submit**.

Uninstall the solution

You can uninstall the Content Localization on AWS solution from the AWS Management Console or by using the AWS Command Line Interface. You must manually delete the Amazon Simple Storage Service (Amazon S3) bucket(s) created by this solution. AWS Solutions Implementations do not automatically delete S3 buckets in case you have stored data to retain.

Using the AWS Management Console

1. Sign in to the [AWS CloudFormation console](#).
2. On the **Stacks** page, select this solution's installation stack.
3. Choose **Delete**.

Using AWS Command Line Interface

Determine whether the AWS Command Line Interface (AWS CLI) is available in your environment. For installation instructions, refer to What Is the AWS Command Line Interface in the *AWS CLI User Guide*. After confirming that the AWS CLI is available, run the following command.

```
$ aws cloudformation delete-stack --stack-name <installation-stack-name>
```

Deleting the Amazon S3 buckets

This solution is configured to retain the solution-created Amazon S3 bucket (for deploying in an opt-in Region) if you decide to delete the AWS CloudFormation stack to prevent accidental data loss. After uninstalling the solution, you can manually delete this S3 bucket if you do not need to retain the data. Follow these steps to delete the Amazon S3 bucket.

1. Sign in to the [Amazon S3 console](#).
2. Choose **Buckets** from the left navigation pane.
3. Locate the <stack-name> S3 buckets.
4. Select the S3 bucket and choose **Delete**.

To delete the S3 bucket using AWS CLI, run the following command:

```
$ aws s3 rb s3://<bucket-name> --force
```

Using the Content Localization on AWS application

Access the web application

After the solution successfully launches, you can access the web application. The solution sends an email containing information to access the web application, including a temporary password. The first time you log in to the application, you will be prompted to change your password.

Identify the URL

Use the following procedure to identify the URL for the web application. This will allow you to sign in.

1. Sign in to the [AWS CloudFormation console](#) and select the solution's stack.
2. On the **Stacks** page, select the **WebStack** nested stack.
3. Choose the **Outputs** tab.
4. Under the **Key** column, locate **ContentLocalizationSolution**, and select the corresponding value.
5. Open the web application in a new tab or browser window.
6. Sign in with your username (Admin email) and temporary password provided in the invitation email.
7. After signing in, follow the prompts to create a new password.

Upload a video and run a workflow

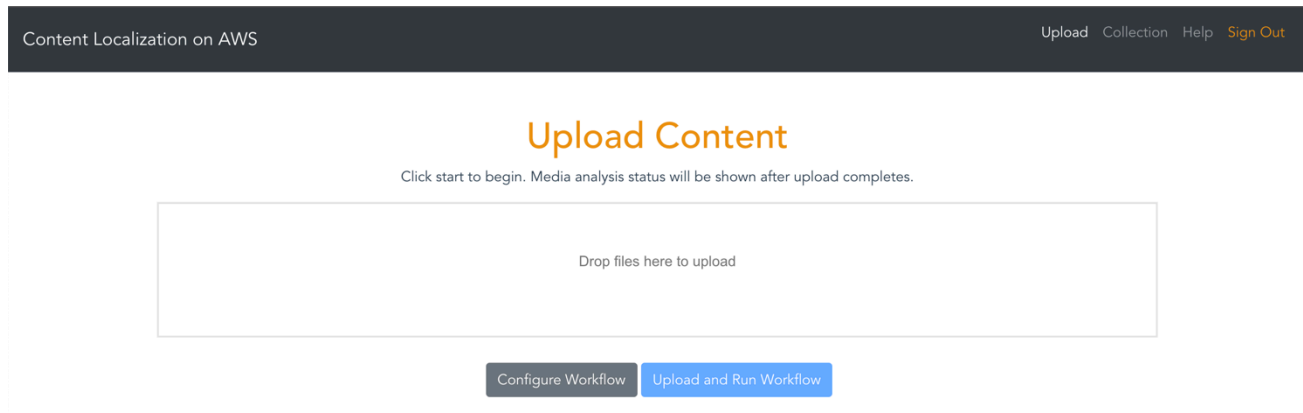
Supported input types

This solution uses AWS Elemental MediaConvert to transcode uploaded videos into the MP4 format required by the analysis operators and therefore supports the same video input video formats. For information about the file formats supported by MediaConvert, refer to [Supported input codecs and containers](#) in the *MediaConvert User Guide*.

Run a workflow

Use the following procedure to upload videos and to start an analysis.

1. Sign in to the Content Localization on AWS web application and choose **Upload**.
2. From the **Upload Content** page, drag and drop one or more media files into the upload box shown in Figure 6.



Content Localization on AWS Upload Content page

3. Choose **Configure Workflow** and select at least one language from the **Target Languages** list.
4. Select a source language for your input from the **Source Language** drop down.
5. (Optional) Select an Amazon Transcribe Custom Vocabulary from the **Custom Vocabulary** drop down.
6. (Optional) Select one or more Amazon Translate Custom Terminologies from the **Custom Terminologies** box.
7. (Optional) Select one or more Amazon Translate Parallel Data sets from the **Parallel Data** box.
8. (Optional) Turn on additional analysis on the input. By default, only the operators required to automatically generate and translate subtitles are turned on. You can optionally select other analysis operators to run in the workflow by selecting the check boxes next to each operator.

Video Operators

- Object Detection
- Technical Cue Detection
- Shot Detection
- Celebrity Recognition
- Face Detection
- Word Detection Face Search

Thumbnail position:

10 sec

Audio Operators

- Transcribe

Source Language

English, US
↕

Custom Vocabulary

(optional)
↕

Custom Language Models

(optional)
↕

Use Existing Subtitles

(optional) Enter .vtt filename

Text Operators

- Comprehend Entities
- Comprehend Key Phrases
- Translate
- Generate audio translations with Amazon Polly

Custom Terminologies: (1 selected)

Perseverance_es_ru (ru,es)

Parallel Data: (1 selected)

SampleParallelData (es,zh)

Target Languages:

Spanish x
Chinese (Simplified) x
Russian x

Type or select

Afrikaans

Albanian

Amharic

Arabic

Armenian

Azerbaijani

Bengali

Bosnian

Bulgarian

Catalan

Chinese (Traditional)

Croatian

Czech

Danish

Dari

Dutch

Estonian

Farsi (Persian)

Filipino (Tagalog)

Finnish

French

French (Canadian)

Georgian

German

Greek

Gujarati

Haitian Creole

Hausa

Hebrew

Hindi

Hungarian

Icelandic

Indonesian

Irish

Italian

Japanese

Kannada

Kazakh

Korean

Latvian

Malay

Malayalam

Maltese

Marathi

Mongolian

Norwegian

Pashto

Polish

Portuguese

Portuguese (Portugal)

Punjabi

Romanian

Serbian

Sinhala

Slovak

Slovenian

Somali

Spanish (Mexico)

Swahili

Swedish

Tamil

Telugu

Thai

Turkish

Ukrainian

Urdu

Uzbek

Vietnamese

Welsh

Content Localization on AWS Operator categories

- To start the analyses, choose **Upload and Run Workflow**. After the workflow has completed, a table appears below the Operator categories to verify that the video was successfully uploaded and analyzed.

Work with the media collection

When you run a workflow on a video through the web application, an Asset ID is created for that video and is used to identify all of the metadata and other outputs created for that video from running workflows. You can browse the collection of media assets in the **Collection** view.

Thumbnail	File Name	Status	Asset ID	Created	Actions
	1-original_how_to_land_on_a_come.tmp4	Started	b3bd8047-0267-468c-b765-4bc9a33a620b	10/1/2021 12:09:56 PM	Analyze Delete
	1_Perseverance.mp4	Started	348b507d-c537-4429-b7e1-98514e7ee441	10/1/2021 12:07:50 PM	Analyze Delete
	sample-video.mp4	Started	7136ca26-a982-41de-9aa4-62dabbccf0ee	10/1/2021 12:04:57 PM	Analyze Delete

Content Localization on AWS Media Collection page

Browse the collection of processed media assets

Sign in to the Content Localization on AWS web application and choose **Collection**.

The **Media Collection** table lists the following attributes for each asset in the media collection:

- **Thumbnail** - a frame sampled from the input video to help visually identify it.
- **File Name** - the name of the input video file.
- **Status** - status of the most recent workflow run on this asset.
- **Asset ID** - the Media Insights on AWS data plane asset ID assigned to this input. This asset ID is used to retrieve all of the metadata and other outputs that were created for this video by running Media Insights on AWS workflows.
- **Created** - the date and time the asset was created.
- **Actions** - buttons to perform further actions on the asset.

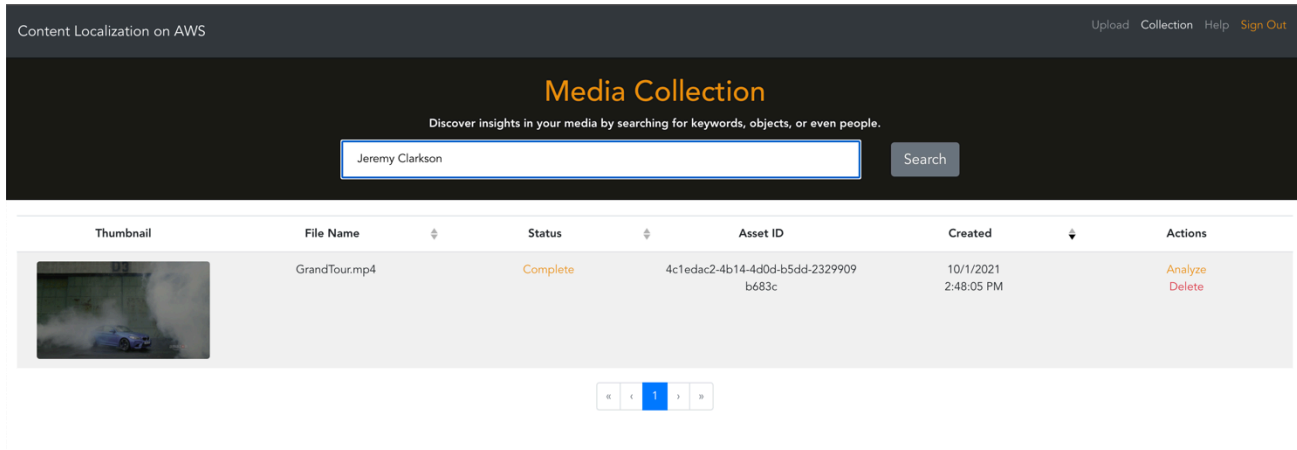
Search the media collection

The **Media Collection** page contains a search bar that can be used to filter the collection to those media assets that contain specific search terms. The workflow operators analyze the media file and indexes and then catalog the information in an Amazon OpenSearch Service instance. You


can search all aspects of a media file using Apache Lucene, the Amazon OpenSearch Service query language. The following examples show common search patterns.

Use full-text queries

Full-text queries allow you to search for any type of data that exists in the video catalog. For example, the Amazon Rekognition celebrity detection service returns the full names of celebrities detected in a video. Figure 9 shows a search for a celebrity by first and last name. Figure 10 shows the search results.



Using the search function in the web application



Search results in the web application

Search high confidence data

After the analysis workflow completes, labels returned by Amazon Rekognition are assigned a confidence value. For details about Amazon Rekognition label detection metadata, refer to [Detecting labels in a video](#). You can use that value to filter search results. For example, `Violence AND Confidence:>80` will search for videos containing violence with an 80% or higher confidence threshold.

Search data from individual operators

Searches query the metadata catalog in Amazon OpenSearch Service. For example, a search for the term *violence* would match videos containing the *violence* label from content moderation and would also match video transcripts that contain the word *violence*. You can restrict your search to focus content moderation results with operator names, for example: `Operator:content_moderation AND (Name:violence AND Confidence:>80)`.

You can use the following operator names to filter search queries:

- `label_detection`
- `celebrity_detection`
- `content_moderation`
- `face_detection`
- `transcribe`
- `key_phrases`
- `entities`
- `webcaptions_<language-code>`

Search related concepts across multiple operators

You can also conduct compound searches using multiple operator names. For example, the following search query returns *violence* identified by content moderation and *guns* or *weapons* identified by label detection: `(Operator:content_moderation AND Name:Violence AND Confidence:>80) OR (Operator:label_detection AND (Name:Gun OR Name:Weapon))`.

View source-language subtitles

The default workflow for this application includes automatically generated WebVTT and SRT subtitles from the Amazon Transcribe output. This part of the workflow can't be deactivated. Subtitles can be edited interactively in the application and saved to invoke reprocessing of downstream operators in the workflow using the updated subtitles as input.

To view subtitles in the application:

1. Run a workflow on a video to create an asset.
2. Choose **Collection**.
3. Locate the media file you want to analyze and under the **Actions** column, choose **Analyze**.
4. From the **Speech Recognition** tab, review the subtitles in the **Subtitles** tab for your content.

When you select a subtitle, the video advances to the location of the subtitle.

The screenshot displays the 'Content Localization on AWS' interface. At the top, there are navigation links for 'Upload', 'Collection', 'Help', and 'Sign Out'. Below this, there are two main tabs: 'Speech Recognition' (selected) and 'ML Vision'. Under 'Speech Recognition', there are sub-tabs for 'Transcript', 'Subtitles', 'Translation', 'KeyPhrases', and 'Entities'. The 'Subtitles' tab is active, showing a list of subtitle entries with their start and end times and the text of the subtitle. The selected subtitle is highlighted in blue. Below the list, there are three buttons: 'Download VTT', 'Save vocabulary', and 'Save edits'. To the right of the list is a video player showing a car on a road. The video player has a progress bar and a play button. Below the video player, there is a section for 'Perform Additional Analysis' with the following details:

- Asset ID: 4c1edac2-4b14-4d0d-b5dd-2329909b683c
- Filename: GrandTour.mp4
- Video duration: 02:18
- Video format: MPEG-4
- Video file size: 228.63 MB
- Video bit rate: 13182293 bps
- Video frame rate: 23.976 fps
- Video resolution: 1920 x 1080
- Audio bit rate: 384 kb/s
- Audio sampling rate: 48.0 kHz
- Encoded date: UTC 2019-09-06 04:50:22
- Source Language: English, US

View subtitles

To download WebVTT and SRT subtitles, choose **Download Subtitles** and select the format you want to download.

Edit source language subtitles

You can edit subtitles content, start time, and end time.


1. After selecting your asset to analyze, from the **Speech Recognition** tab, choose the **Subtitles** tab.
2. Select a subtitle. When you select a subtitle, the video advances to the location of the subtitle.
3. Select the form box you want to edit for a specific subtitle (start time, end time, or subtitle content) and enter the new value.
4. To save the changes you made to the subtitles, choose **Save Edits** at the bottom of the page. The solution saves the new subtitles to the Media Insights on AWS data plane (in WebCaptions format) and reprocesses any downstream operators for this asset that take the subtitles (WebCaptions) as an input.

Use corrections to source language subtitles to create an Amazon Transcribe custom vocabulary

You can use the corrections you make while editing subtitles to generate an Amazon Transcribe custom vocabulary that you can use in future workflows to improve the quality of Amazon Transcribe results for your content.

1. After making edits in the **Subtitles** tab for your content, choose **Save Custom Vocabulary** at the bottom of the page.
2. (Optional) To add to an existing vocabulary, use the radio button to select it before editing. This populates the vocabulary table in the form with the content of the existing vocabulary plus any edits you made to the subtitles.
3. The form presents you with a table that is pre-populated with corrections that you made to the subtitles for the asset. You can further modify the rows in the table by selecting a cell and entering new values. You can also add and delete rows using the (+) and (-) buttons at the end of each row. The table contains the following values for each row:
 - **Original phrase** - the phrase that was in the automatically generated subtitles. This is provided for reference and can't be edited in the form.
 - **New phrase** - the phrase that replaced the original phrase through the editor.

- **Sounds Like** (optional) - The pronunciation of your word or phrase using the standard orthography of the language to mimic the way that the word sounds. For details, refer to [Custom vocabularies](#) in the *Amazon Transcribe Developer Guide*.
- **IPA** (optional) - The pronunciation of your word or phrase using IPA characters. For details, refer to [Custom vocabularies](#) in the *Amazon Transcribe Developer Guide*.
- **Display As** - Defines how you want the word or phrase looks when it's output. For details, refer to [Custom vocabularies](#) in the *Amazon Transcribe Developer Guide*.



Save Vocabulary? dialog box

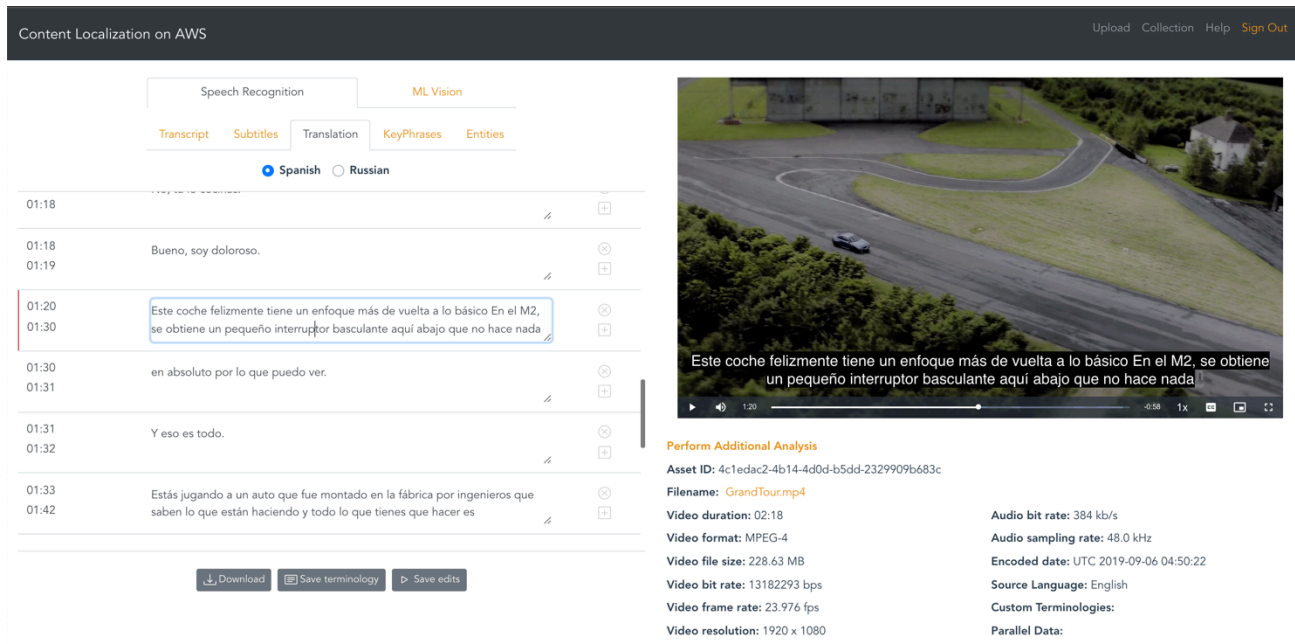
4. If you are creating a new vocabulary, fill in a name for the vocabulary in the **Vocabulary Name** field.
5. Choose **Save**.

View target-language subtitles

The default workflow for this application automatically generates WebVTT and SRT subtitles for all target languages specified by the application. This part of the workflow can't be deactivated. Subtitles can be edited interactively in the application and saved to invoke reprocessing of downstream operators in the workflow using the updated subtitles as input.

1. Run a workflow on a video to create an asset.
2. Choose **Collection**.

3. Locate the media file you want to analyze and under the **Actions** column, choose **Analyze**.
4. From the **Speech Recognition** tab, review the translated subtitles in the **Translation** tab for your content.
5. Select the language you want to view. When you select a subtitle, the video advances to the location of the subtitle.



The screenshot shows the 'Content Localization on AWS' interface. At the top, there are tabs for 'Speech Recognition' and 'ML Vision'. Under 'Speech Recognition', there are sub-tabs for 'Transcript', 'Subtitles', 'Translation', 'KeyPhrases', and 'Entities'. The 'Translation' tab is active, and the language is set to 'Spanish'. A list of subtitles is shown on the left, with the following text highlighted in blue:

Este coche felizmente tiene un enfoque más de vuelta a lo básico En el M2, se obtiene un pequeño interruptor basculante aquí abajo que no hace nada

On the right, a video player shows the same subtitle overlaid on a video of a car on a track. Below the video player, there are technical details for the asset:

- Perform Additional Analysis
- Asset ID: 4c1edac2-4b14-4d0d-b5dd-2329909b683c
- Filename: GrandTour.mp4
- Video duration: 02:18
- Video format: MPEG-4
- Video file size: 228.63 MB
- Video bit rate: 13182293 bps
- Video frame rate: 23.976 fps
- Video resolution: 1920 x 1080
- Audio bit rate: 384 kb/s
- Audio sampling rate: 48.0 kHz
- Encoded date: UTC 2019-09-06 04:50:22
- Source Language: English
- Custom Terminologies:
- Parallel Data:

Target-language translation

Download target-language WebVTT and SRT subtitles

1. After analyzing your asset, choose the **Translation** tab for your content.
2. Select the language you want to work with.
3. Choose **Download** and select the format you want to download.

Download target-language audio file

The content localization workflow creates a synthesized audio output using Amazon Polly for each target language. This audio output is intended to be used as audio-only and, therefore, does not preserve the timing with the video content. If Amazon Polly does not support a target-language, then the processing is skipped and the download button will be inactive for that language.

1. After analyzing your asset, choose the **Translation** tab for your content.
2. Select the language you want to work with.
3. Choose **Download** and select **Download Audio** from the dropdown.

Edit target language subtitles

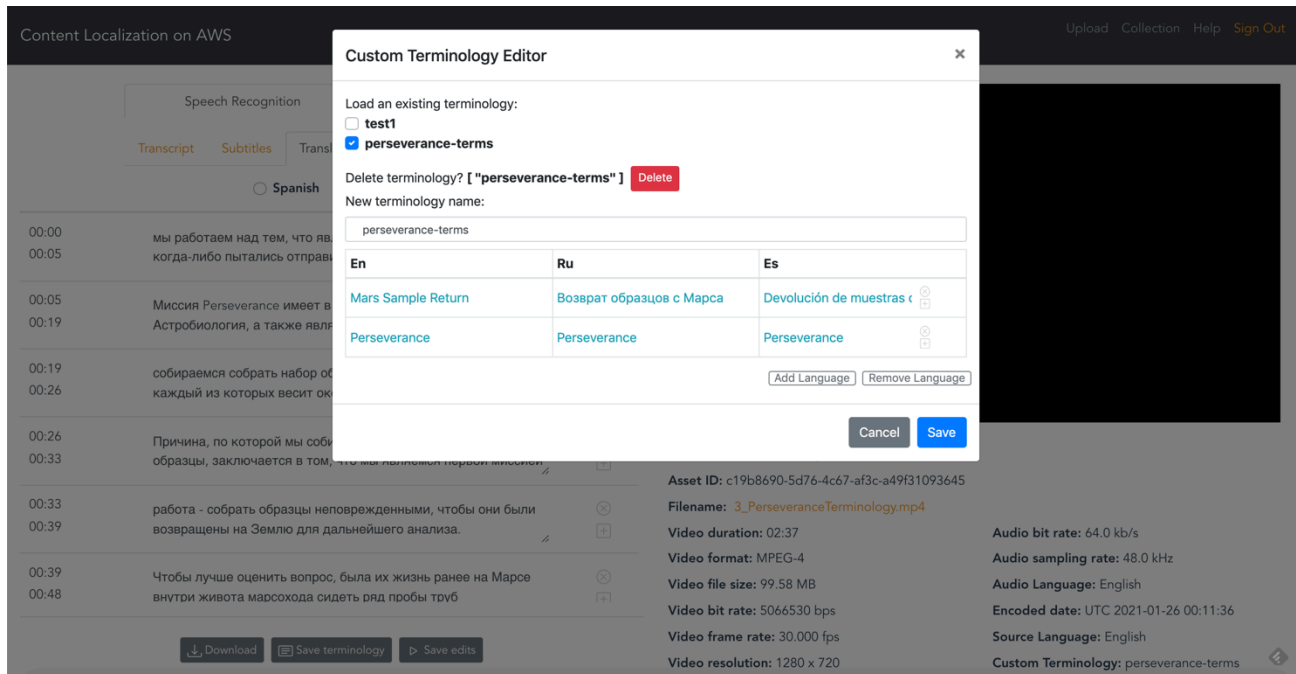
1. After analyzing your asset, choose the **Translation** tab for your content.
2. Select the target language you want to work with.
3. Select a subtitle. When you select a subtitle, the video advances to the location of the subtitle.
4. Select the form box you want to edit for a specific subtitle (start time, end time or subtitle content) and enter the new value.
5. To save the changes you made to the subtitles, choose **Save Edits** at the bottom of the page. The solution saves the new subtitles to the Media Insights on AWS data plane (in WebCaptions format) and reprocesses any downstream operators for this asset that take the translated subtitles (WebCaptions_<target-language>) as an input.

Use corrections to target language subtitles to create an Amazon Translate terminology

You can use the corrections you make while editing target-language subtitles to generate an Amazon Translate terminology that you can use in future workflows to customize the Amazon Translate results for your content.

1. After making edits in the **Subtitles** tab for your content, choose **Save Terminology** at the bottom of the page.
2. (Optional) To add to an existing terminology, use the radio button to select it before editing. You can combine existing terminologies by selecting multiple radio buttons. This populates the vocabulary table in the form with the content of the existing terminologies plus any edits you made to the subtitles.
3. The form presents you with a table that is pre-populated with corrections you have made to the subtitles for the asset. You can further modify the rows in the table by selecting a cell and entering new values. Add and delete rows using the (+) and (-) buttons at the end of each row. Add and delete languages in the form using the **Add Language** and **Remove Language** buttons.

- The terminology table contains columns for the source language and for each target language specified in the workflow. Any target-language edits for a phrase will be filled in in the table. You must fill in all of the cells for other languages you want to include in the terminology or remove the language from the table.
- Enter the name you want to use for this terminology. If you use the name of an existing terminology, it will be replaced with the contents from this form.
- Choose **Save**.



Custom Terminology Editor dialog box

Create user for the application

If more than one person needs access to the web application, the solution administrator can set up additional users in Amazon Cognito.

- Sign in to the [Amazon Cognito console](#).
- Choose **Manage User Pools**.
- In the **Your User Pools** page, select the name of the user pool containing the prefix **MI**.
- On the **MieUserPool** page, from the left navigation pane, choose **Users and Groups**.
- On the **Users** tab, choose **Create user**.

6. In the **Create user** dialog box:
 - a. Enter a username.
 - b. Enter a temporary password. Verify the options to send an invitation to the user and the verifications for phone number and email are not selected.
 - c. Choose **Create user**.
7. On the **MieUserPool** page, under the **Username** column, select the user you just created.

The screenshot shows the 'Users' tab for a user pool named 'MieUserPool-Oq0I23Q4CoqC'. The 'Users' tab is active, and the 'Create user' button is visible. Below the buttons is a table with the following columns: Username, Enabled, Account status, Email verified, Phone number verified, Updated, and Created. The table contains two rows of user data. The second row, for the user 'myuser@example.com', is circled in red.

Username	Enabled	Account status	Email verified	Phone number verified	Updated	Created
[blurred]	Enabled	CONFIRMED	-	-	Jun 24, 2020 4:12:26 PM	Jun 24, 2020 3:48:43 PM
myuser@example.com	Enabled	FORCE_CHANGE_PASSWORD	-	-	Jun 24, 2020 4:35:43 PM	Jun 24, 2020 4:35:43 PM

Amazon Cognito user pools page

8. On the **Users** page, choose **Add to group**.
9. In the **Add user** dialog box, access the drop-down list and select **MieDevelopersGroup**.

The screenshot shows the 'Add user to group' dialog box. The title is 'Add user myuser@example.com to group'. Below the title is a search box containing 'MieDevelopersGroup' with a red arrow pointing to it. At the bottom of the dialog, the 'Add to group' button is circled in red.

Amazon Cognito Add user to group dialog box

The user can now access the web application, upload media files, and run the analysis workflows.

Developer guide

Source code

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

Reference

This section includes information about an optional feature for collecting unique metrics for this solution, pointers to related resources, and a list of builders who contributed to this solution.

Anonymized data collection

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

- **Solution ID** - The AWS solution identifier
- **Unique ID (UUID)** - Randomly generated, unique identifier for each Content Localization on AWS deployment
- **Timestamp** - Data-collection timestamp
- **Version** - The version of the solution that was deployed
- **CFTemplate** - The CloudFormation action that activated the metrics report

AWS owns the data gathered through this survey. Data collection is subject to the [AWS Privacy Notice](#). To opt out of this feature, complete the following steps before launching the AWS CloudFormation template.

1. Download the [AWS CloudFormation template](#) to your local hard drive.
2. Open the AWS CloudFormation template with a text editor.
3. Modify the AWS CloudFormation template mapping section from:

```
AnonymizedData:  
  SendAnonymizedData:  
    Data: Yes
```

to:

```
AnonymizedData:  
  SendAnonymizedData:  
    Data: No
```

4. Sign in to the [AWS CloudFormation console](#).
5. Select **Create stack**.
6. On the **Create stack** page **Specify template** section, select **Upload a template file**.
7. Under **Upload a template file**, choose **Choose file** and select the edited template from your local drive.
8. Choose **Next** and follow the steps in [Launch the stack](#) in the Deploy the solution section of this guide.

To opt out of data collection in the nested Media Insights on AWS stack, use one of the following options:

Option 1: Deploy Media Insights on AWS separately with metric reporting disabled

1. Follow the instructions to disable metric reporting in the [Anonymized data collection](#) section of the *Media Insights on AWS implementation guide*, then deploy the Media Insights on AWS solution.
2. Follow the instructions to [install Content Localization on AWS over an existing Media Insights on AWS stack](#).

Option 2: Modify the source template used to deploy Media Insights on AWS as a nested stack

1. Download the Media Insights on AWS template from <https://solutions-reference.s3.amazonaws.com/media-insights-on-aws/<version>/media-insights-on-aws-stack.template> and replace <version> with the Media Insights on AWS version you would like to deploy.
2. Follow the instructions to disable metric reporting in the [Anonymized data collection](#) section of the *Media Insights on AWS implementation guide*, but do not deploy the Media Insights on AWS template.
3. Upload the template to an S3 bucket.
4. Edit the [Content Localization on AWS template](#) and change TemplateURL within the MieStack properties to point to the template you've uploaded to S3.

From

```
MieStack:
  Type: "AWS::CloudFormation::Stack"
  Properties:
    TemplateURL: !Join
      - ""
      - - "https://solutions-reference.s3.amazonaws.com/media-insights-on-aws/"
        - !FindInMap
          - MediaInsights
          - Release
          - Version
        - "/media-insights-on-aws-stack.template"
```

To

```
MieStack:
  Type: "AWS::CloudFormation::Stack"
  Properties:
    TemplateURL: "<new url>"
```

5. Use the instructions in the README to [build](#) and [deploy](#) Content Localization on AWS.

Related resources

Related AWS Solutions

- The [Media Insights on AWS](#) solution is a dependency.

Contributors

- Alexandria Burkleaux
- Brandon Dold
- Ian Downard
- Jill Adams
- Tom Gilman
- San Dim Ciin

- Colin McCoy
- Eric Thoman
- David Chung

Revisions

Date	Change
November 2021	Initial release
February 2022	Version 2.0.0: Add support for using custom language models with Amazon Transcribe. For more information, refer to the CHANGELOG.md file in the GitHub repository.
August 2022	Version 2.0.1: Updated the Python version in the Lambda functions from version 3.6 to version 3.8. For more information, refer to the CHANGELOG.md file in the GitHub repository.
January 2023	Version 2.0.2: Updated Python libraries for enhanced security. For more information, refer to the CHANGELOG.md file in the GitHub repository.
April 2023	Version 2.1.0: Updated Media Insight Engine references to Media Insights on AWS. Added a Service Catalog AppRegistry resource to register the CloudFormation template and underlying resources as an application in both AWS Service Catalog AppRegistry and AWS Systems Manager Application Manager. Mitigated impact caused by new default settings for S3 Object Ownership (ACLs disabled) for all new S3 buckets. Updated implementation guide to new documentation template. For more information, refer to the CHANGELOG.md file in the GitHub repository.
July 2023	Version 2.1.1: Upgraded to Node.js to 18 and JS SDK to v3. Updated parameter names for

Date	Change
	consistency. For more information, refer to the CHANGELOG.md file in the GitHub repository.
October 2023	Version 2.1.2: Updated solution to use latest Media Insights on AWS functionality. Added options for opting out of data collection in the nested Media Insights on AWS stack. Bug fixes. For more information, refer to the CHANGELOG.md file in the GitHub repository.
November 2023	Documentation update: Added Confirm cost tags associated with the solution to the Monitoring the solution with AWS Service Catalog AppRegistry section.
March 2024	Version 2.1.3: Updated solution to use latest Media Insights on AWS version. Upgraded vulnerable packages. Added packages and minor frontend changes for the Vue 2 to 3 migration build. Enabled versioning on website bucket. For more information, refer to the CHANGELOG.md file in the GitHub repository.
March 2024	Version 2.1.4: Minor package update. For more information, refer to the CHANGELOG.md file in the GitHub repository.
June 2024	Version 2.1.5: Upgraded vulnerable packages. For more information, refer to the CHANGELOG.md file in the GitHub repository.
July 2024	Version 2.1.6: Upgraded vulnerable packages. For more information, refer to the CHANGELOG.md file in the GitHub repository.

Date	Change
August 2024	Version 2.1.7: Updated package versions to resolve security vulnerabilities. Updated solution to use latest Media Insights on AWS functionality. For more information, refer to the CHANGELOG.md file in the GitHub repository.
September 2024	Version 2.1.8: Upgraded vulnerable packages. For more information, refer to the CHANGELOG.md file in the GitHub repository.
October 2024	Version 2.1.9: Upgraded vulnerable packages. For more information, refer to the CHANGELOG.md file in the GitHub repository.
October 2024	Version 2.1.10: Upgraded vulnerable packages. For more information, refer to the CHANGELOG.md file in the GitHub repository.
December 2024	Version 2.1.11: Upgraded vulnerable packages. For more information, refer to the CHANGELOG.md file in the GitHub repository.

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.

Content Localization on AWS is licensed under the terms of the [Apache License Version 2.0](#).