

The background features a dark blue gradient with abstract geometric shapes. On the left, a thin orange line forms a triangle. On the right, a large curved shape transitions from blue to orange. The text is centered in the upper right area.

AWS re:Invent

NOV. 29 – DEC. 3, 2021 | LAS VEGAS, NV

MDS 305

Automate video clipping and highlight generation

Aravindharaj Rajendran
Solutions Developer
AWS

Kyle Huang
Data Scientist
AWS



Agenda

Use cases and challenges

Automate video clipping – AI/ML techniques, available AWS services and solutions

Introduction – Media Replay Engine (MRE)

Demo – Automated clipping of football (soccer) game using MRE

Q&A

Terminology

Detection

Identify useful **data** from video that brings **context** and establishes **priority**

Segmentation

Identify interesting **segments** or **clips** in video using the **data** detected

Highlight

Calculate the best subset of **clips** that meet specific criteria based on the **data** detected

Use cases

Sports



Summarized re-broadcast
Catchup highlights
Post-game highlights
Personalized highlights
Social media feeds

News & entertainment



Generate clips for news shows
Awards shows
Shopping networks

Enterprise



Index keynotes
Training content
Video logging

Customer challenges

Reduce cost to edit live content




Growing amount of content sources combined with a multitude of replay channel needs without adding significant operational expense with production

Enable personalized highlights for DTC



Producing engaging content for consumers doesn't scale with manual editing methods

Establish new revenue through digital channels



Insertion of personalized ads to target audiences is complex to orchestrate

AI/ML techniques

Segmentation

Scene classification
(camera view changes)

Audio-based detection

Audio classification (detect applause, whistle, ball hit)

Speech detection (specific phrases)

Vision-based detection

Proximity between two objects
(player to net)

Pose estimation (trick shots, handshake, interesting action)

Object detection (trophy, scoreboard, logo)

Face detection (specific players)

AWS services and solutions

Key AWS services

AWS Elemental MediaLive

Amazon Rekognition

Amazon SageMaker

Amazon Transcribe

AWS Lambda

AWS Step Functions

AWS solutions

Media2Cloud

Media Insights Engine (MIE)

Media Replay Engine (MRE)

Media Replay Engine (MRE)



MRE – What is it?



An **open-source** AWS framework you deploy into your own AWS account



Automatically process **live and VOD** broadcast content to identify segments (clips) of interest using machine learning






Leverage sample **plugins** and ML model **notebooks** to process your content



Schedule **replays** for linear and DTC channels to enable **personalized** content and revenue generation through ad insertion

How MRE works



MRE solution framework

A media processing specific pipeline that is customized with plugins and ML models you bring

Add plugins and ML models

Leverage community-provided plugins or build your own for your video/audio content

Configure content specific profiles

Configure reusable profiles per content type that apply the desired plugins and ML models to live video

Connect your video sources

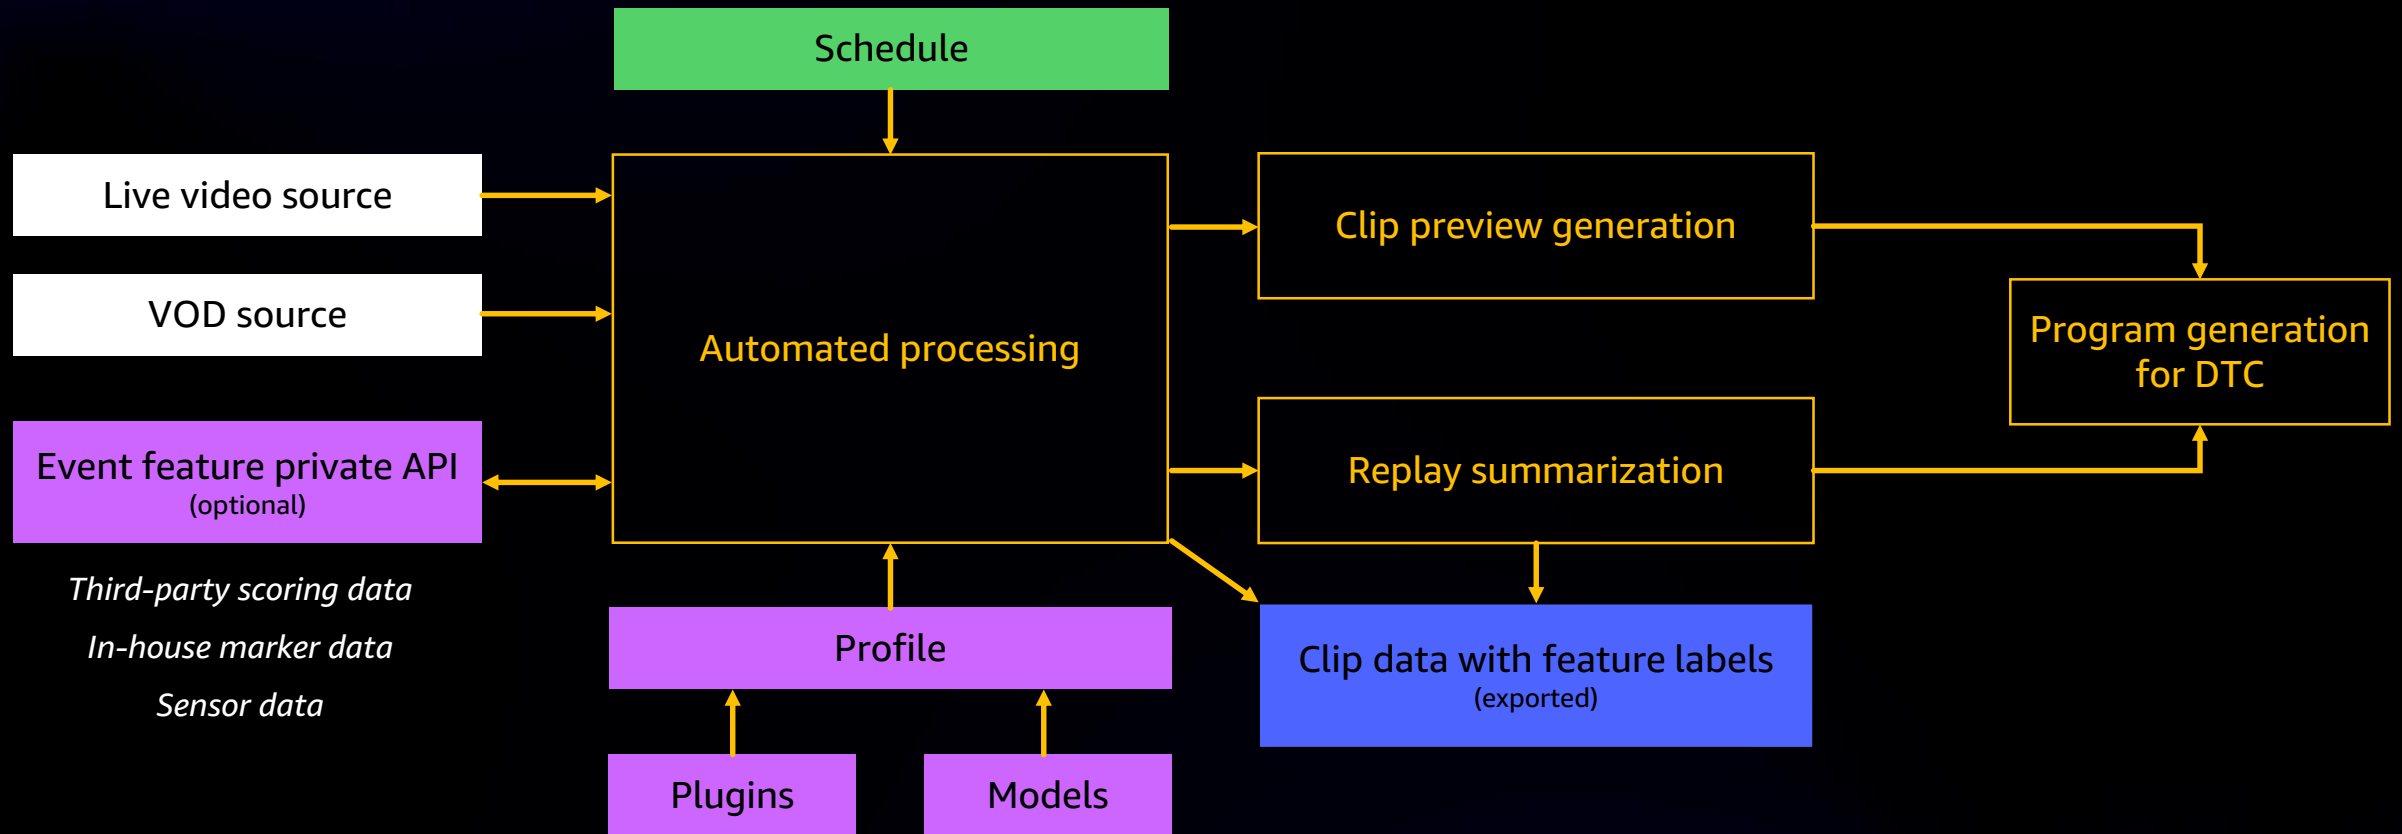
Configure channel endpoints to stream video to, and enable video with, your broadcast sources

Schedule events and replays

Add events to be processed on the media channels, apply the selected profile, and process replay requests live or after the event

MRE system diagram

Admin Producer Editor Developer




MRE provided components

MRE Core
(framework and automation)



Automation



Optimized pipeline

Secure API



MRE Console UI*
(to manage MRE)



Event List

Event Time	Program Name	Event	Channel	Duration (seconds)	Status	Export	Actions
July 19th 2021, 1:00:00 pm	ML Game	Portland vs Seattle	MLN	120	In Progress		⋮
July 17th 2021, 8:00:00 pm	Football game	Portland vs Seattle	MLN	120	Completed		⋮
July 19th 2021, 8:17:00 pm	Football game	Goal	MLN	1	Completed		⋮
July 19th 2021, 8:17:00 pm	Football game	Goal	MLN	1	In Progress		⋮
July 19th 2021, 8:00:00 pm	Football game	Goal	MLN	100	Completed		⋮
July 19th 2021, 8:00:00 pm	Football game	Goal	MLN	100	Completed		⋮

Event registry

Clips Detected

Thumbnail	Start Time (seconds)	Label	Length (seconds)
	104.455	TBD	35.033
	126.628	TBD	7.655
	97.899	TBD	16.514
	67.659	TBD	25.823

Segment (clip) auditing




Clip reviewing

*Optional




MRE Samples*
(plugins, models, fan UX)



Plugin classes



AI/ML/CV based plugins



Fan experience UI



*Optional



Football (soccer) video clipping demo

ML classification model

- Determine what type of scene it is (contextual classification)
- Sample key frames to establish common patterns
- Pattern identification will help identify where the action is



Segmentation


- Apply AI pattern matching
- Identify start and end of desired segment
- Use provided helper library to store results




or



then



then




or


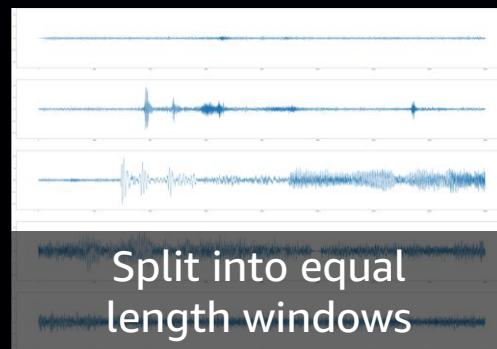




Pose detection

- ML model to output keypoints for each human figure
- Keypoint coordinates can be used with additional plugins for pose detection



Audio scene detection



Starter questions

- *What type of live content have you experimented with segmenting using automation?*
- *What techniques have you tried (CV, ML, API, etc.)?*
- *What challenges have you encountered bringing a solution like this to production scale?*

What's next?



<https://github.com/awslabs/aws-media-replay-engine>

Thank you!

Aravindharaj Rajendran

araviraj@amazon.com

Kyle Huang

kylehh@amazon.com

