
AWS
re:Invent
DECEMBER 2 – 6, 2024 | LAS VEGAS, NV

MAE301

Building resilient live video streaming with AWS Media Services

Kiran Patel
(he/him)
Senior Product Manager
AWS (AWS Elemental)

Bryan Samis
(he/him)
Principal Solutions Architect
AWS

Jamie Mullan
(he/him)
Senior Solution Architect
AWS

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

01 One to many to chaos

02 Levels of resilience

03 Types of failure

04 Service SLA vs. workflow uptime

05 Standard vs. single pipeline

06 Cross-Region workflows

07 Media Quality Confidence Score (MQCS)

08 Trade-offs and considerations

aws

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

One to many to chaos: The OG

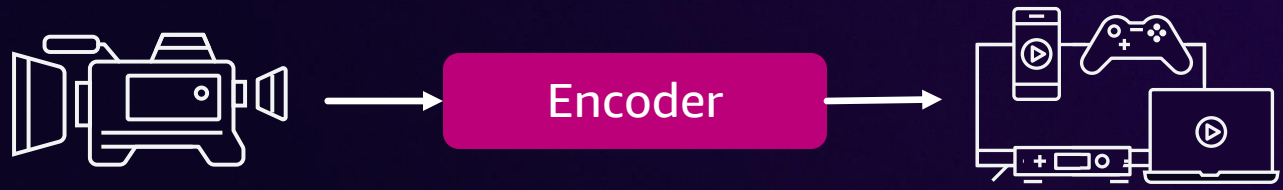

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

One to many to chaos: The breakdown

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

One to many to chaos: Double it up

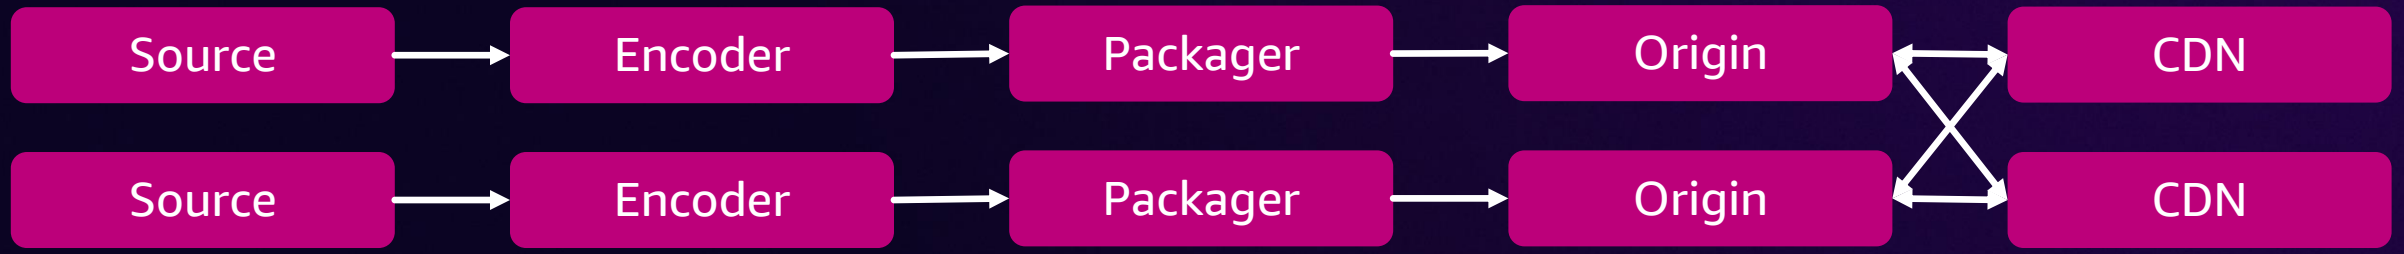

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

One to many to chaos: More is better…?

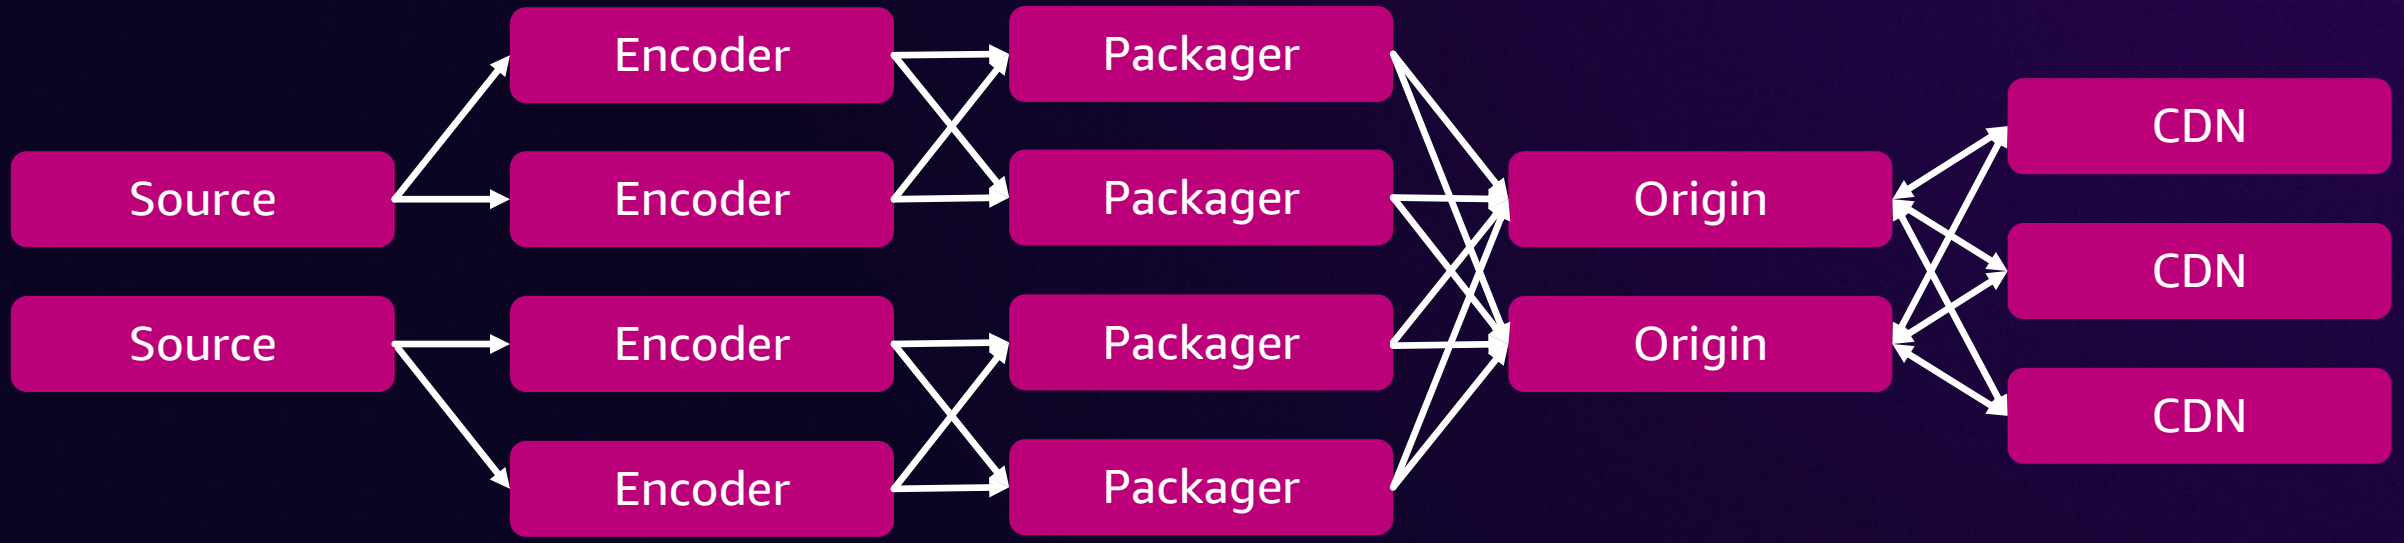

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

One to many to chaos: More is complex

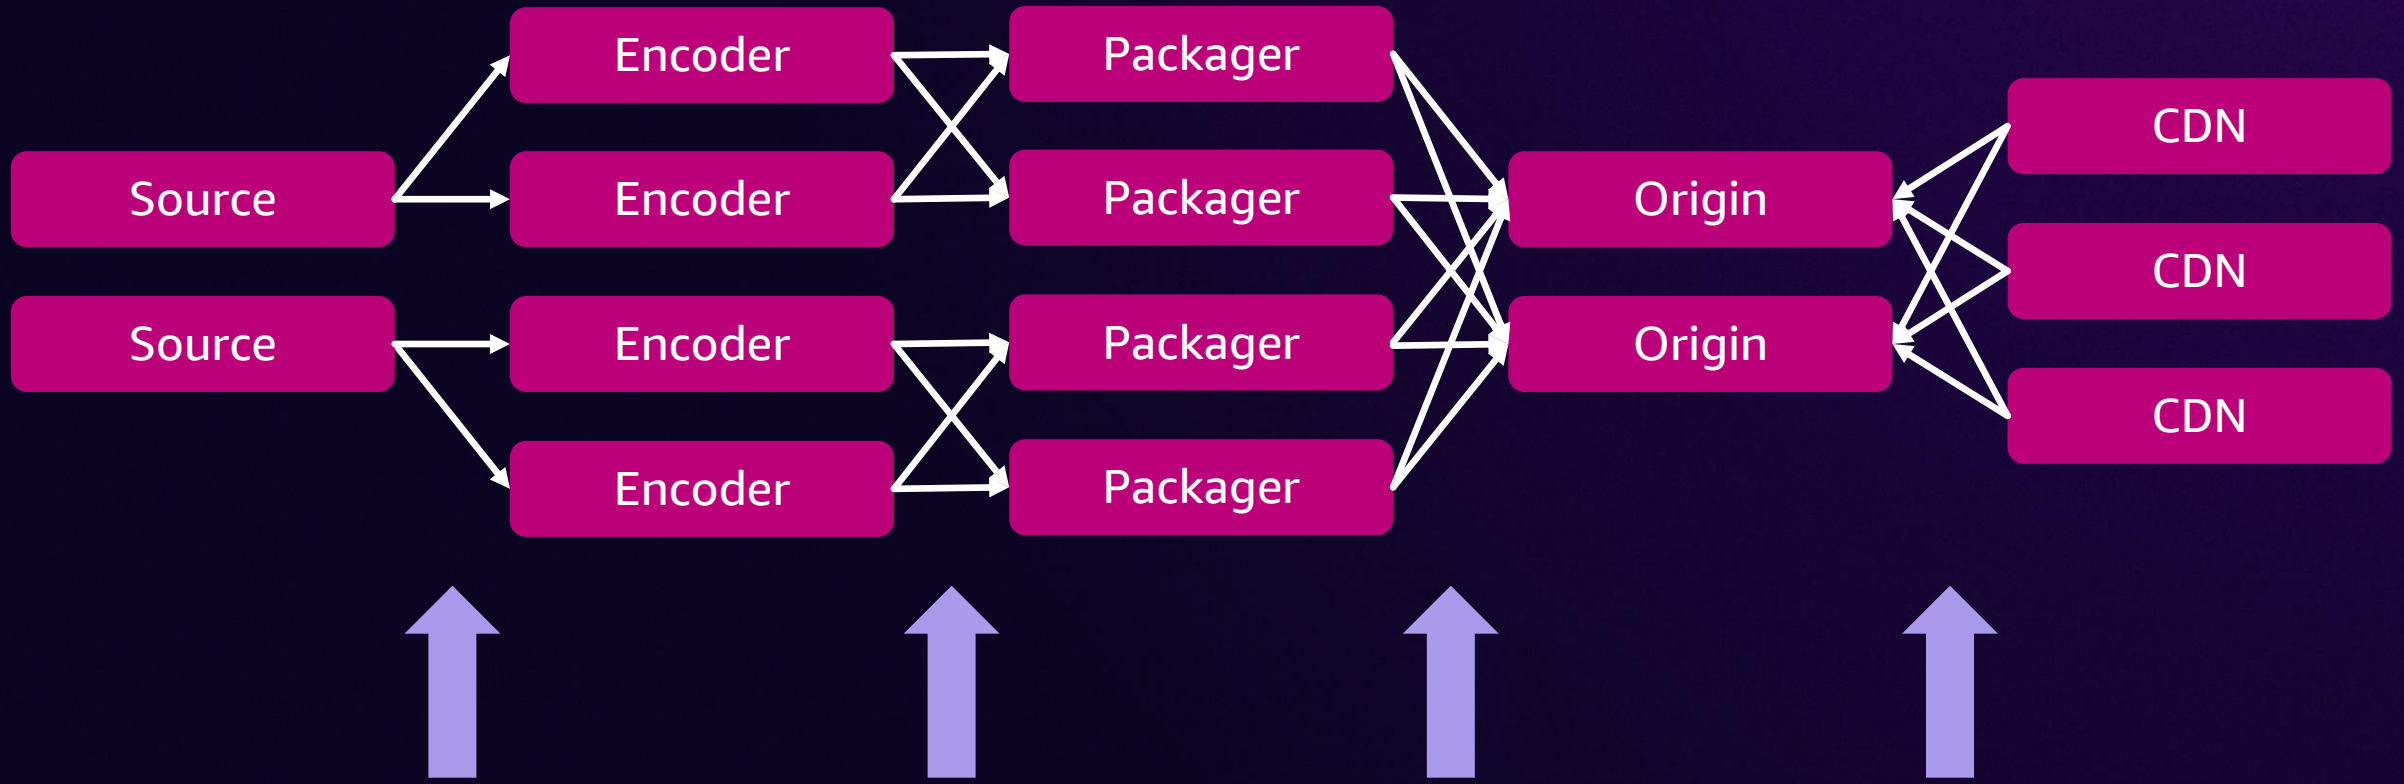

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Levels of resilience

Resilience =

Redundancy

+ Failover

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Levels of resilience: Simple

Simple resilience =

Redundancy by duplicating components

+ Manual failover

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Levels of resilience: Better

Better resilience =

Cloud-based redundancy with auto scaling and healing

+ Automatic failover

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Types of failure

Total outage

Total network outage
Video loss

Partial outage

Partial network outage
Video degraded

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Media Services

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Acceptable downtime with percentage uptime target			
	1 day	30 day month	365 day year
95%	1.2 hours	36 hours	438 hours
	72 minutes	2160 minutes	26280 minutes
	4320 seconds	129600 seconds	1576800 seconds
99%	0.24 hours	7.2 hours	87.6 hours
	14.4 minutes	432 minutes	5256 minutes
	864 seconds	25920 seconds	315360 seconds
99.9%	0.02 hours	0.72 hours	8.76 hours
	1.44 minutes	43.2 minutes	525.6 minutes
	86.4 seconds	2592 seconds	31536 seconds
99.95%	0.01 hours	0.36 hours	4.38 hours
	0.72 minutes	21.6 minutes	262.8 minutes
	43.2 seconds	1296 seconds	15768 seconds
99.99%	0 hours	0.07 hours	0.88 hours
	0.14 minutes	4.32 minutes	52.56 minutes
	8.64 seconds	259.2 seconds	3153.6 seconds
99.999%	0 hours	0.01 hours	0.09 hours
	0.01 minutes	0.43 minutes	5.26 minutes
	0.86 seconds	25.92 seconds	315.36 seconds

aws

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Single-pipeline channels

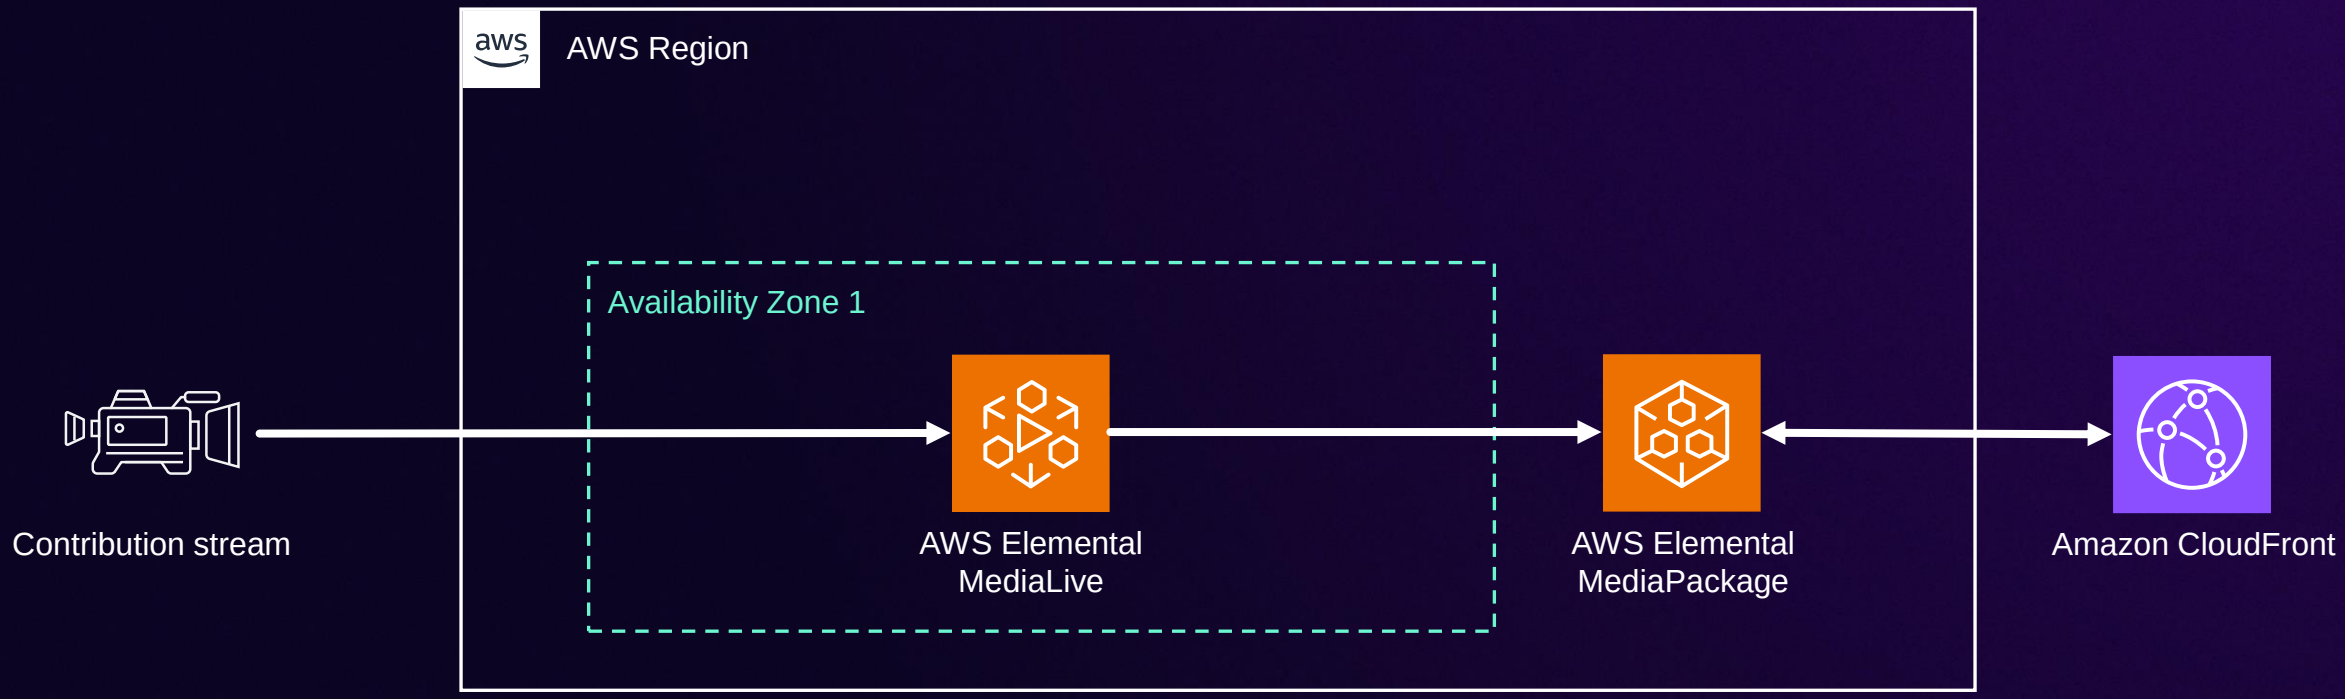

- Lowest cost workflow
- Single threaded – susceptible to individual issues in the source, encoder instance, Availability Zone, or Region

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Standard channels

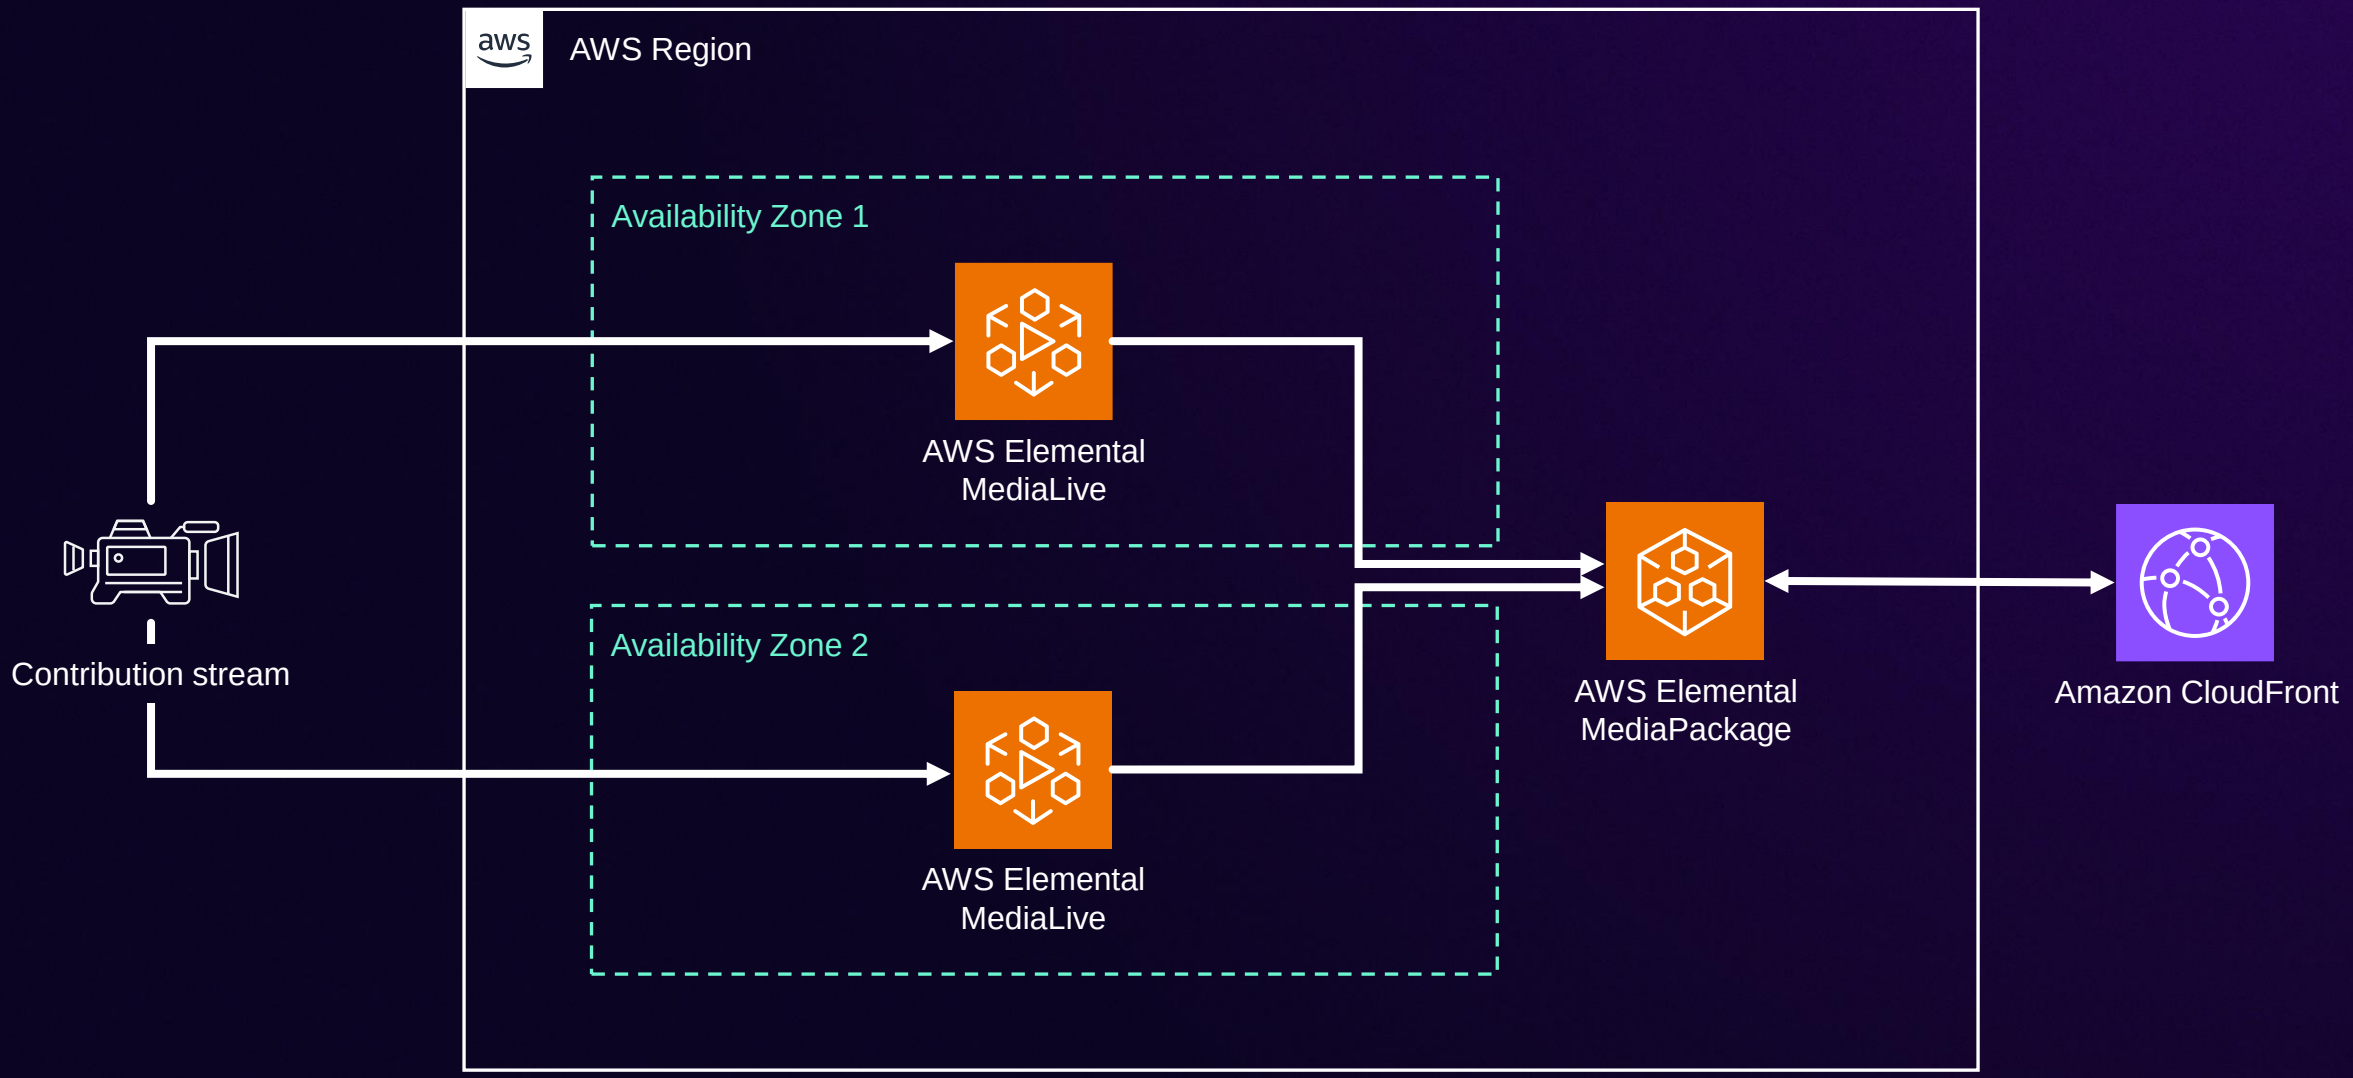

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Work backwards from your SLA

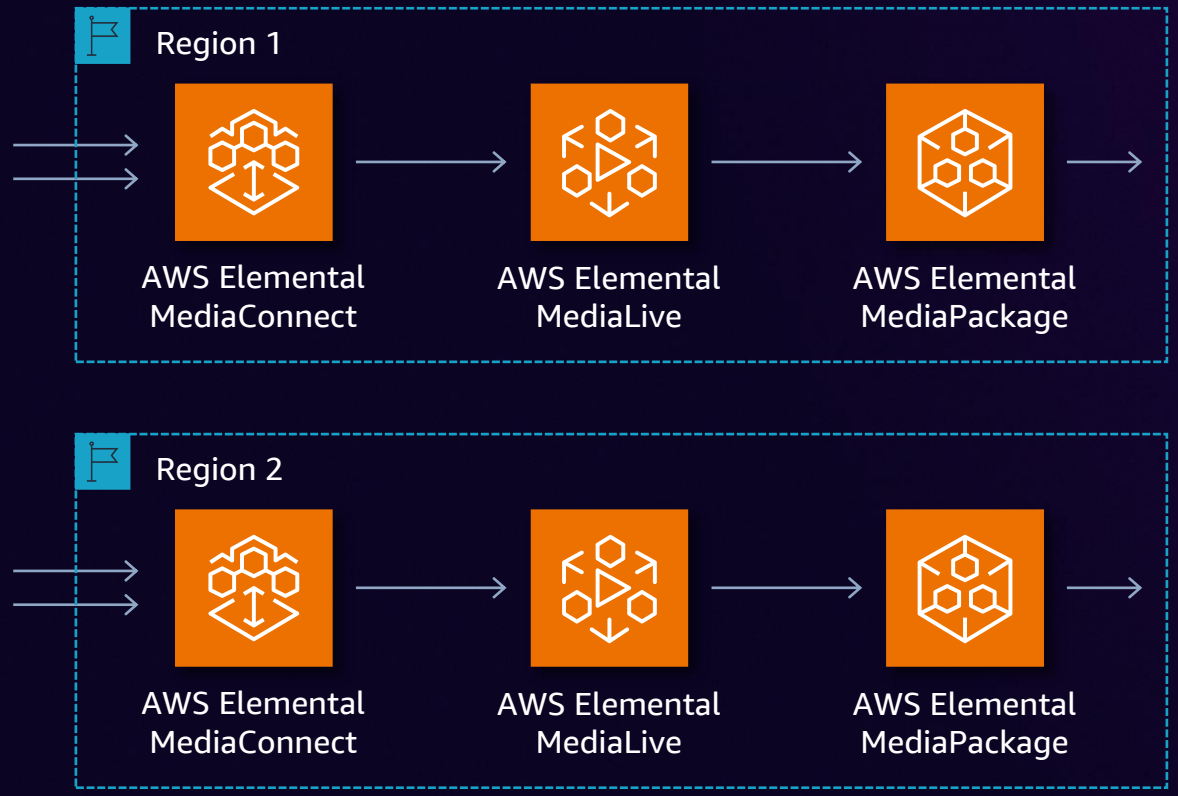

MediaConnect, MediaLive, and MediaPackage each have an SLA of 99.9%

Availability = Availability of x ∗ Availability of y

$(A=A_xA_y)$ etc.

99.9% * 99.9% * 99.9% = 99.7%

$$A=1-((1-A_x)(1-A_y))$$

$A=1-(1-99.7\%)^2$ = 99.9991%

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Seamless cross-Region failover

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Transparent cross-Region failover workflow

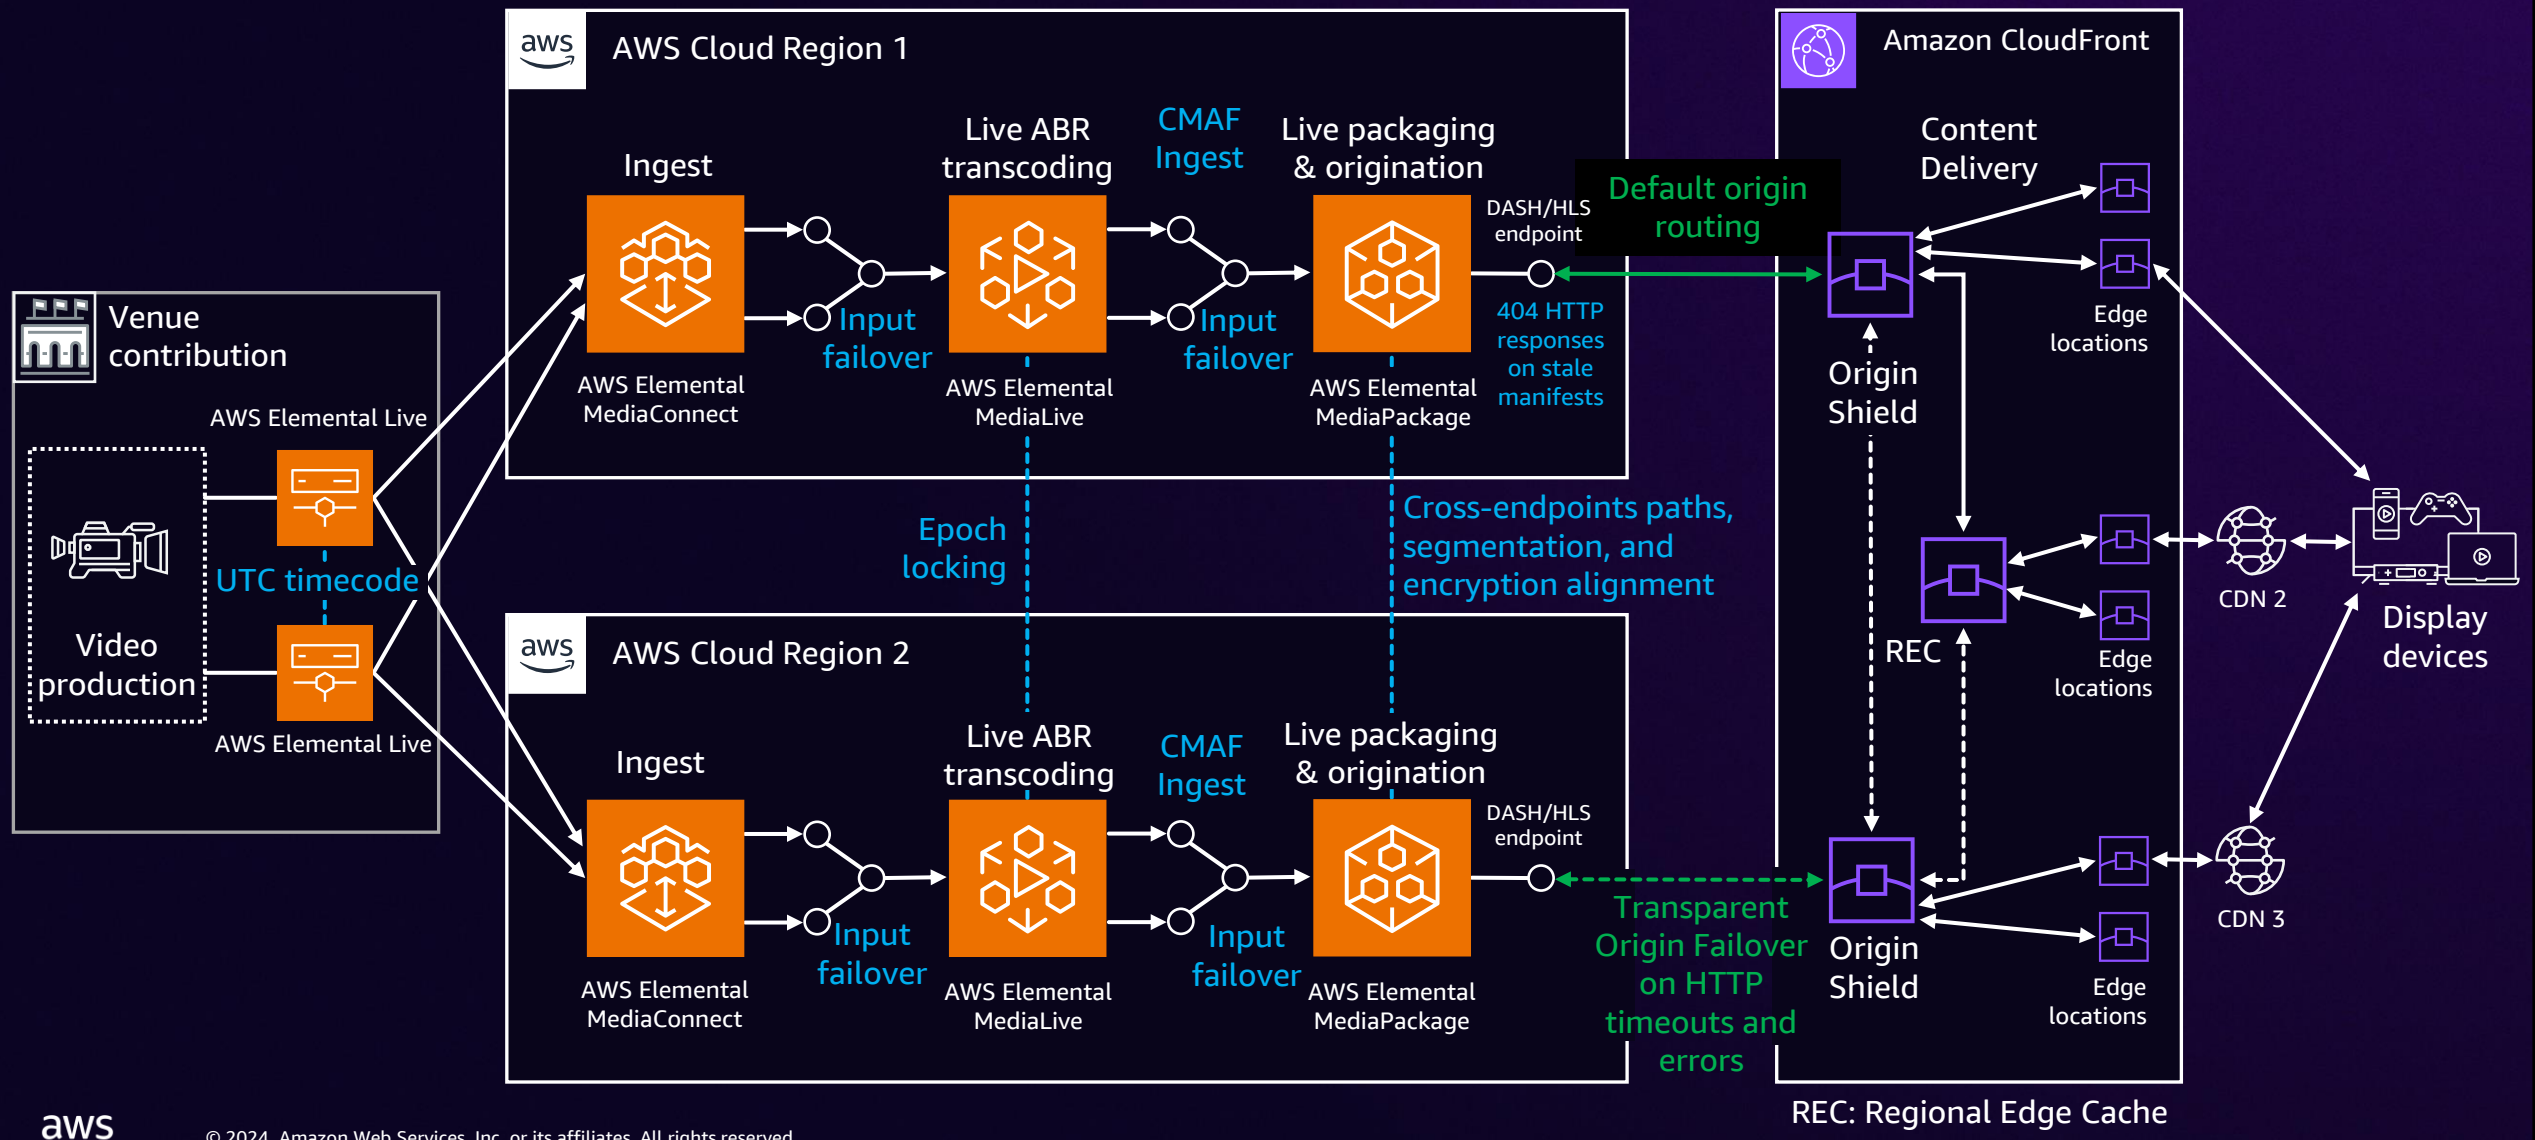

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

CMAF Ingest and epoch timecode locking

- To be “seamless,” all the redundant legs require a frame-accurate alignment of streams across multiple origins
- CMAF Ingest (Interface-1) provides a modern ingest approach between ABR encoders and packagers, leveraging CMAF encapsulation
  - Common Media Application Format
- Regular segmentation cadence (even with SCTE markers) anchored on Unix epoch time
  - Allows predictable packaging and key rotations
  - Enables stateless synchronization of the packagers, based on the source timecode
- Cross-Region origin failover from Content Delivery Network (CDN)
  - Failover parameters: Stale manifests, slate input, incomplete manifests, missing DRM keys

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Seamless cross-region switching between Ohio & Oregon

Learn more about cross-Region failover

AWS for M&E Blog

Build a resilient cross-region live streaming architecture on AWS

by Jamie Mullan, Andrew Fayle, and Christer Whitehorn | on 09 AUG 2024
| in Amazon CloudFront, AWS Elemental MediaLive, AWS Elemental MediaPackage, Direct-to-Consumer & Streaming, Industries, Media & Entertainment, Media Services, Networking & Content Delivery
| Permalink | Comments | Share

Introduction

Customers with live streaming video use cases often express a desire for flexibility in how they deploy live streaming services. Not all live channels are alike, and the ability to balance resiliency against cost is crucial. AWS Media Services from Amazon Web Services (AWS) are designed with redundancy in mind. Stateful services for video transport and encoding, such as AWS Elemental MediaConnect and AWS Elemental MediaLive, can be deployed in multiple availability zones within an AWS Region. Stateless services for video transcoding, origination, and distribution, including AWS Elemental MediaConvert, AWS Elemental MediaPackage, and AWS Elemental MediaTailor, natively span availability zones without any additional user configuration.

Until now, there hasn't been an easy option for customers to fail a live streaming service seamlessly from one region to another. This capability is particularly important for high-profile, high-value live streaming events. This blog post explains how this is now possible and what is involved in doing so.

An example of automatic seamless switching

Fig 1: Seamless switching between regions

Resources

AWS for M&E
AWS Media Services
AWS Resources for M&E
How-To Guides for M&E
AWS Media Blog Home

Follow

Twitter
LinkedIn
YouTube
Email Updates
AWS Elemental LinkedIn

AWS Events
Discover the latest AWS events in your region
Learn more »

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

What about quality-aware failover?

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Media quality-aware resiliency workflow

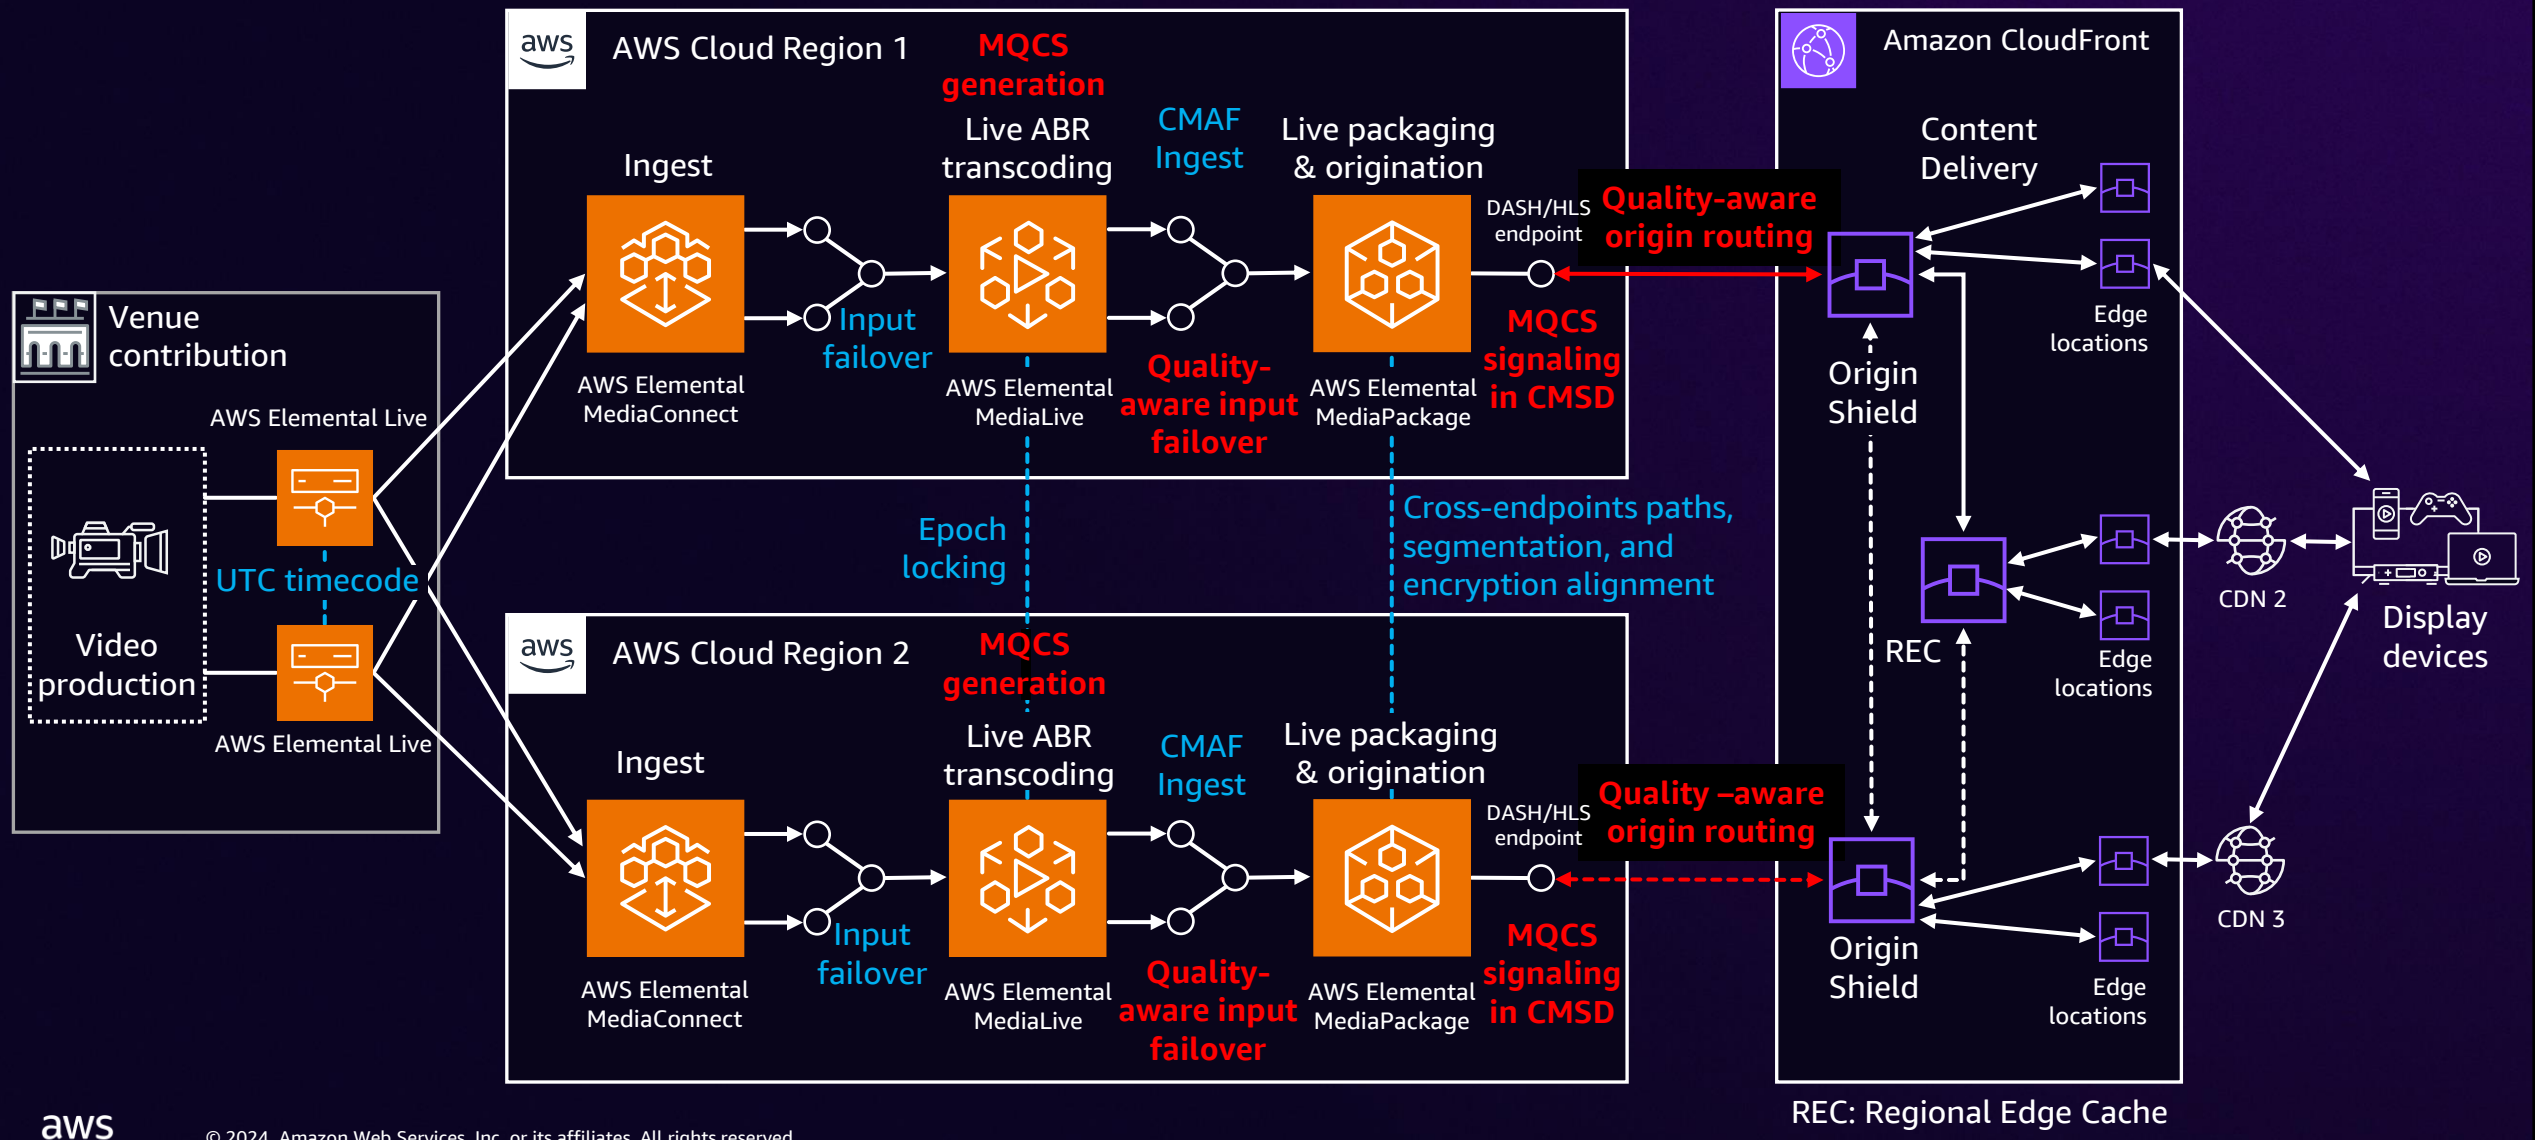

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Media quality-aware resiliency

- AWS Elemental MediaLive generates an Media Quality Confidence Score (MQCS) for each segment, based on multiple input parameters
  - Such as: Input source loss, black frame detection, and frozen frame detection with more parameters coming soon
- AWS Elemental MediaPackage provides in-Region quality-aware failover
- Amazon CloudFront provides cross-Region quality-aware origin selection

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Learn more about media quality-aware resiliency

AWS for M&E Blog

Improve your viewers' live streaming experience with Media-Quality Aware Resiliency

by Tal Shalom and Nicolas Weil | on 21 NOV 2024 | in Amazon CloudFront, Announcements, AWS Elemental MediaLive, AWS Elemental MediaPackage, Direct-to-Consumer & Streaming, Industries, Media & Entertainment, Media Services, Networking & Content Delivery | Permalink | Comments | Share

Architecting a highly resilient solution for streaming premium live sports events, concerts, or news is critical to delight viewers with a high quality of experience (QoE). Deploying your video delivery workflow in two different Amazon Web Services (AWS) Regions is one way to provide extra redundancy. Region failover can mitigate impacts to workflows and even withstand the unlikely event of an partial or complete region failure.

An additional reason to switch between redundant video streams is when a quality degradation is detected in the live video originated from the venue to one of the regions. Traditionally, such quality degradation is detected by a dedicated video expert who maintains "eyes-on-glass" (also known as confidence monitoring) to catch quality issues and switch to a secondary stream. This manual process can take several minutes, during which viewers might miss a climactic moment in a sporting event or an important news item. The longer it takes to recover, the higher the impact to QoE which can eventually lead dissatisfied viewers to cancel subscriptions.

Today we launched Media Quality-Aware Resiliency (MQAR), an integrated capability between Amazon CloudFront and AWS Elemental Media Services. MQAR provides automated, cross-region origin selection and failover based on a dynamically computed media quality score. With MQAR, your "eyes-on-glass" is always on, detecting video quality issues and switching between redundant video streams automatically in matter of seconds.

Always on "eyes-on-glass" for high quality of experience

With AWS, live event streaming starts at the venue with video delivered onto the cloud for processing and distribution to

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

How much resilience do you “need”?

- What does this mean for levels of redundancy?
- What does this mean for ways to fail over?
- How long can you accept running at reduced resilience?
- Can you accept glitches, buffering, or any “dead” airtime?

- What trade-offs are you willing to make?
  - Higher cost
  - Increases end-to-end latency
  - More complex workflow to monitor and maintain

aws

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

Please complete the session survey in the mobile app

Kiran Patel
Linkedin

Bryan Samis
Linkedin

Jamie Mullan
Linkedin

aws

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.