



**ARC303**

# **Pure Play Video, Over-The-Top**

**A Microservices Architecture in the Cloud**

Alex Smith, ASEAN Media Solutions Architect, AWS  
Vidhya Narayanan, Director, Verizon onCue

October 2015



# The Problem


# Your Consumers Have Changed

More content choices

More devices

More delivery methods

New experiences




# Your Business Has Changed

More data

Faster go to market

Leaner teams

Innovation and failure





# What to Expect from the Session

~~Content is king~~

Building from scratch

Battle-hardened lessons – Verizon onCue

# Clarifications

OTT, OVP, AVOD, TVOD, SVOD

It doesn't matter

Microservices

Buzzword?

Specific scope, interoperable services

Allow rapid innovation



# Over the Top Platform



**Content  
Production**

**Content  
Storage**





```
graph LR; A[Content Production] --> B[Content Storage]; B --> C[Processing & Management];
```

**Content  
Production**

**Content  
Storage**

**Processing &  
Management**



```
graph LR; A[Content Production] --> B[Content Storage]; B --> C[Processing & Management]; C --> D[Content Distribution]
```

**Content  
Production**

**Content  
Storage**

**Processing &  
Management**

**Content  
Distribution**




**Content  
Production**

**Content  
Storage**

**Processing &  
Management**

**Content  
Distribution**

**Content  
Consumption**





# Your Existing Physical Requirements

HD-SDI

Satellite transmission

Local SAN

HW transcoder




# Your OTT Approach

FTP / Accelerator

Large NAS

DAM / Workflow System



# Storage and Ingest — Serverless





# Storage and Ingest — Serverless



Amazon Glacier

Content provider



*Lifecycle policies*



S3 IA

*Lifecycle policies*




AWS Import/ Export



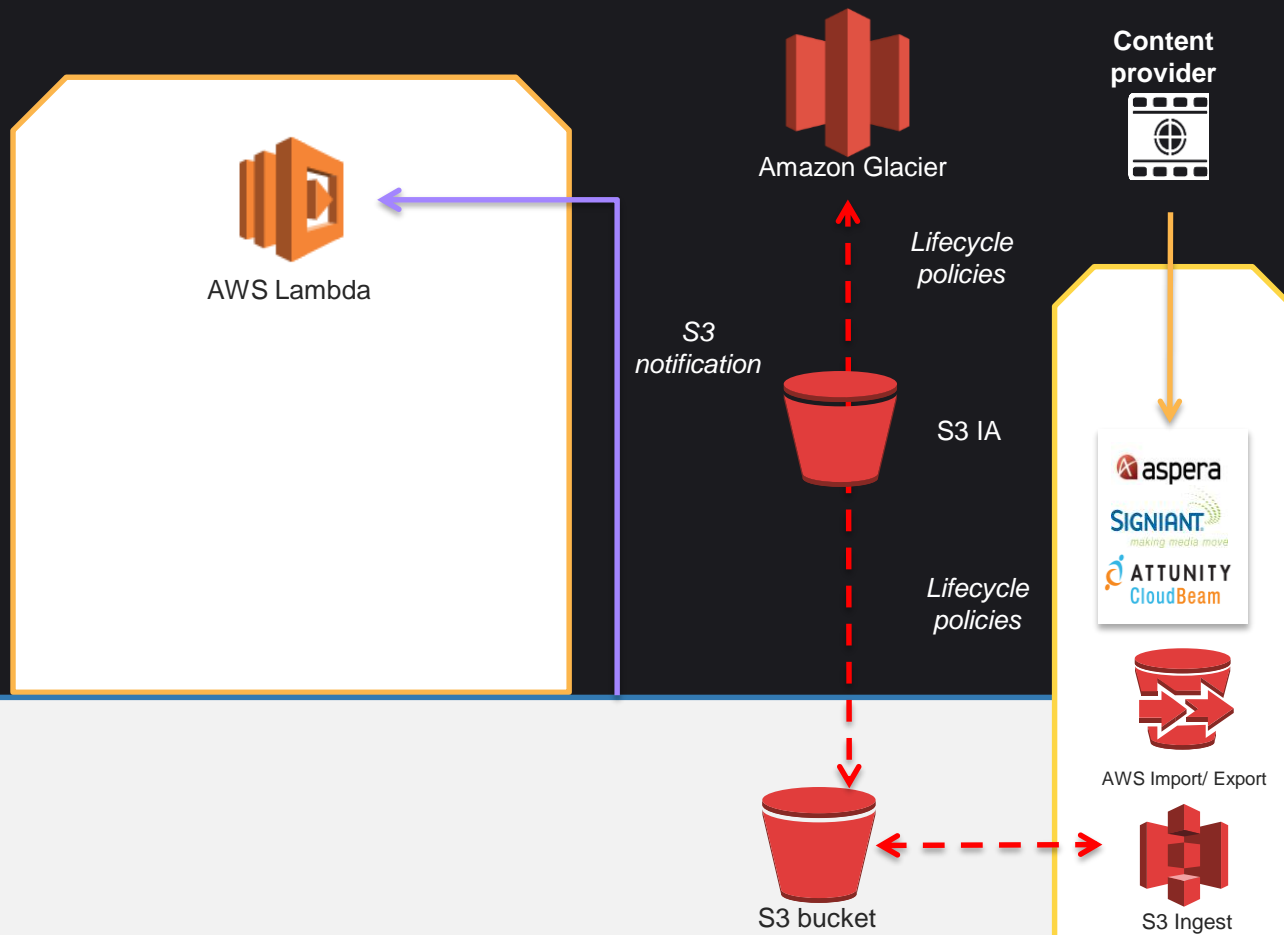
S3 bucket



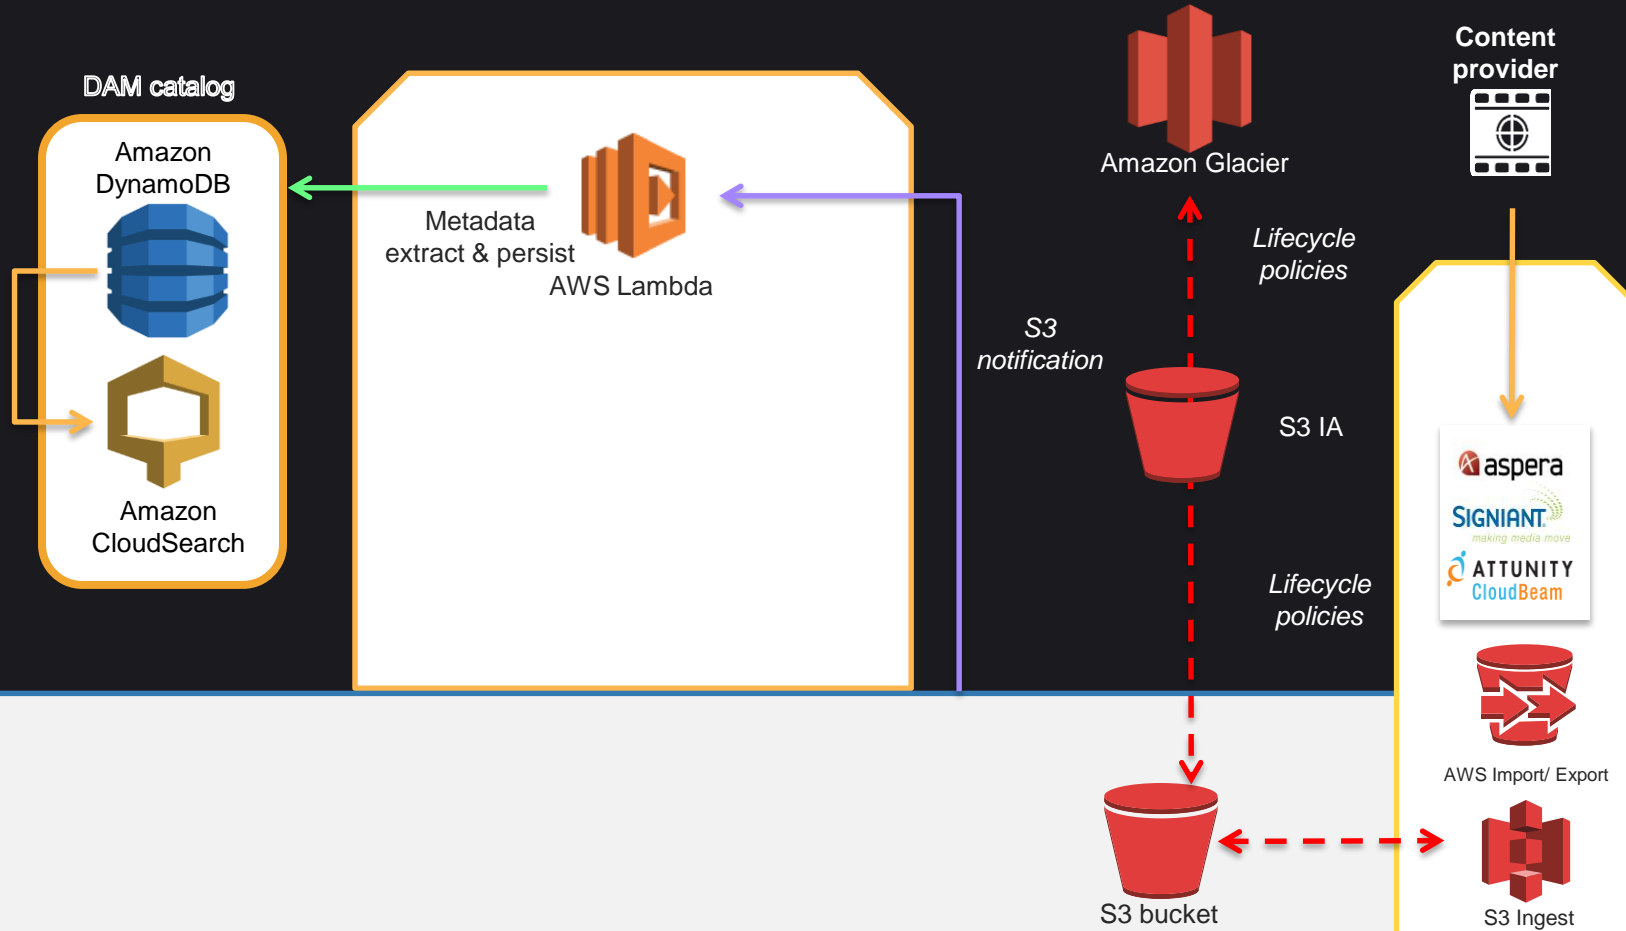
S3 Ingest



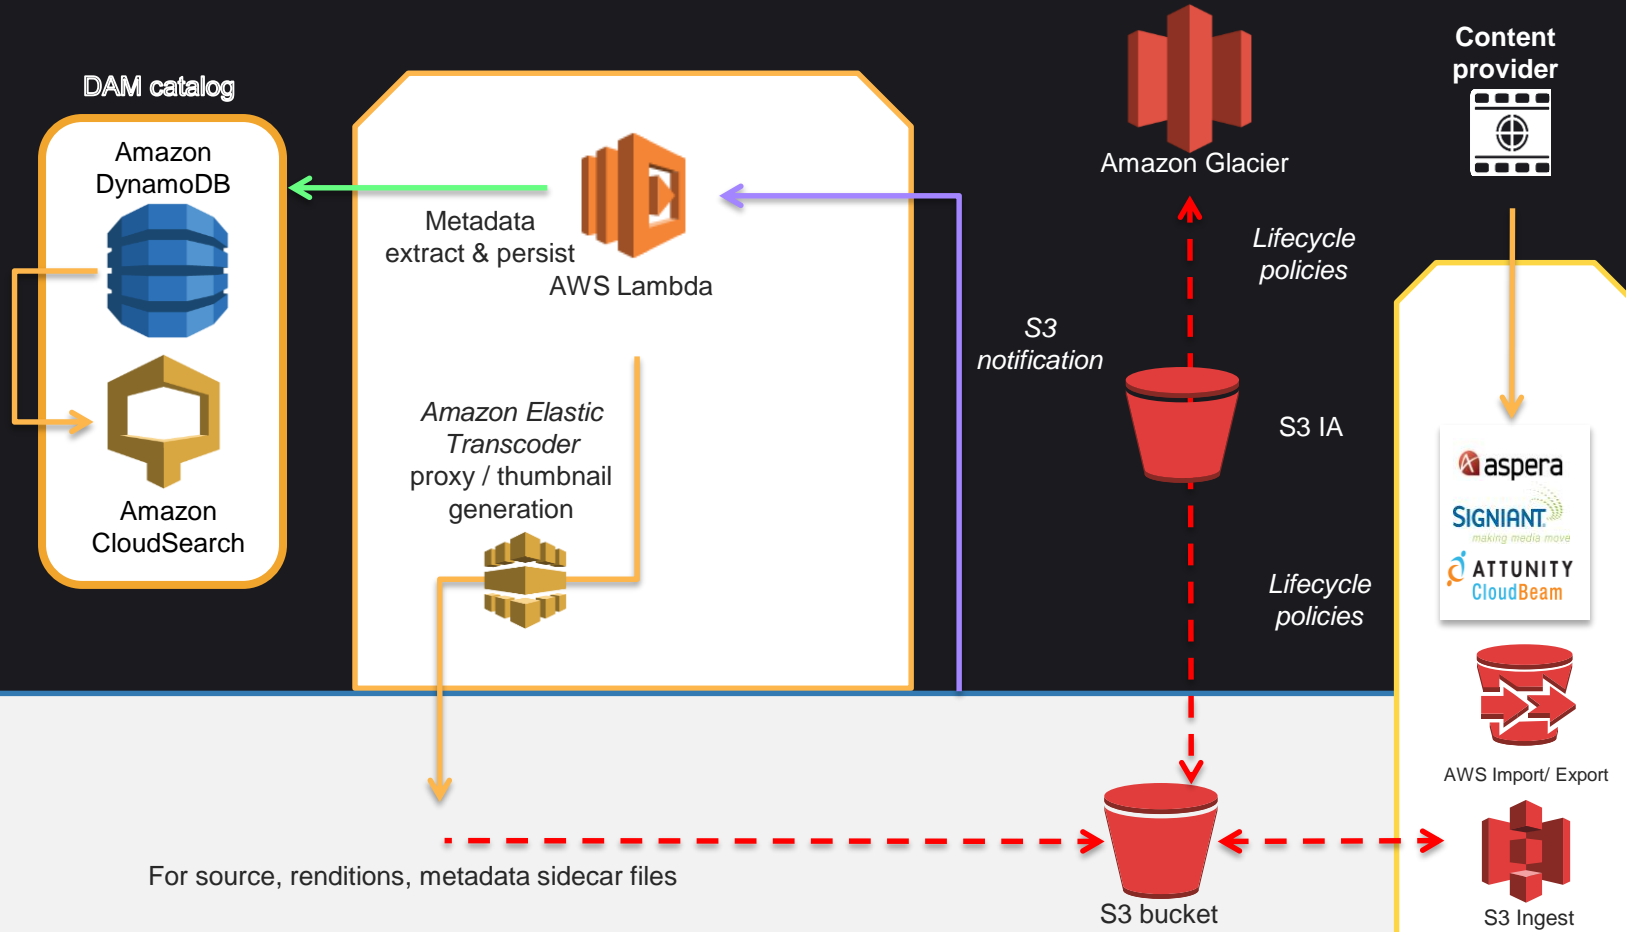
# Storage and Ingest — Serverless



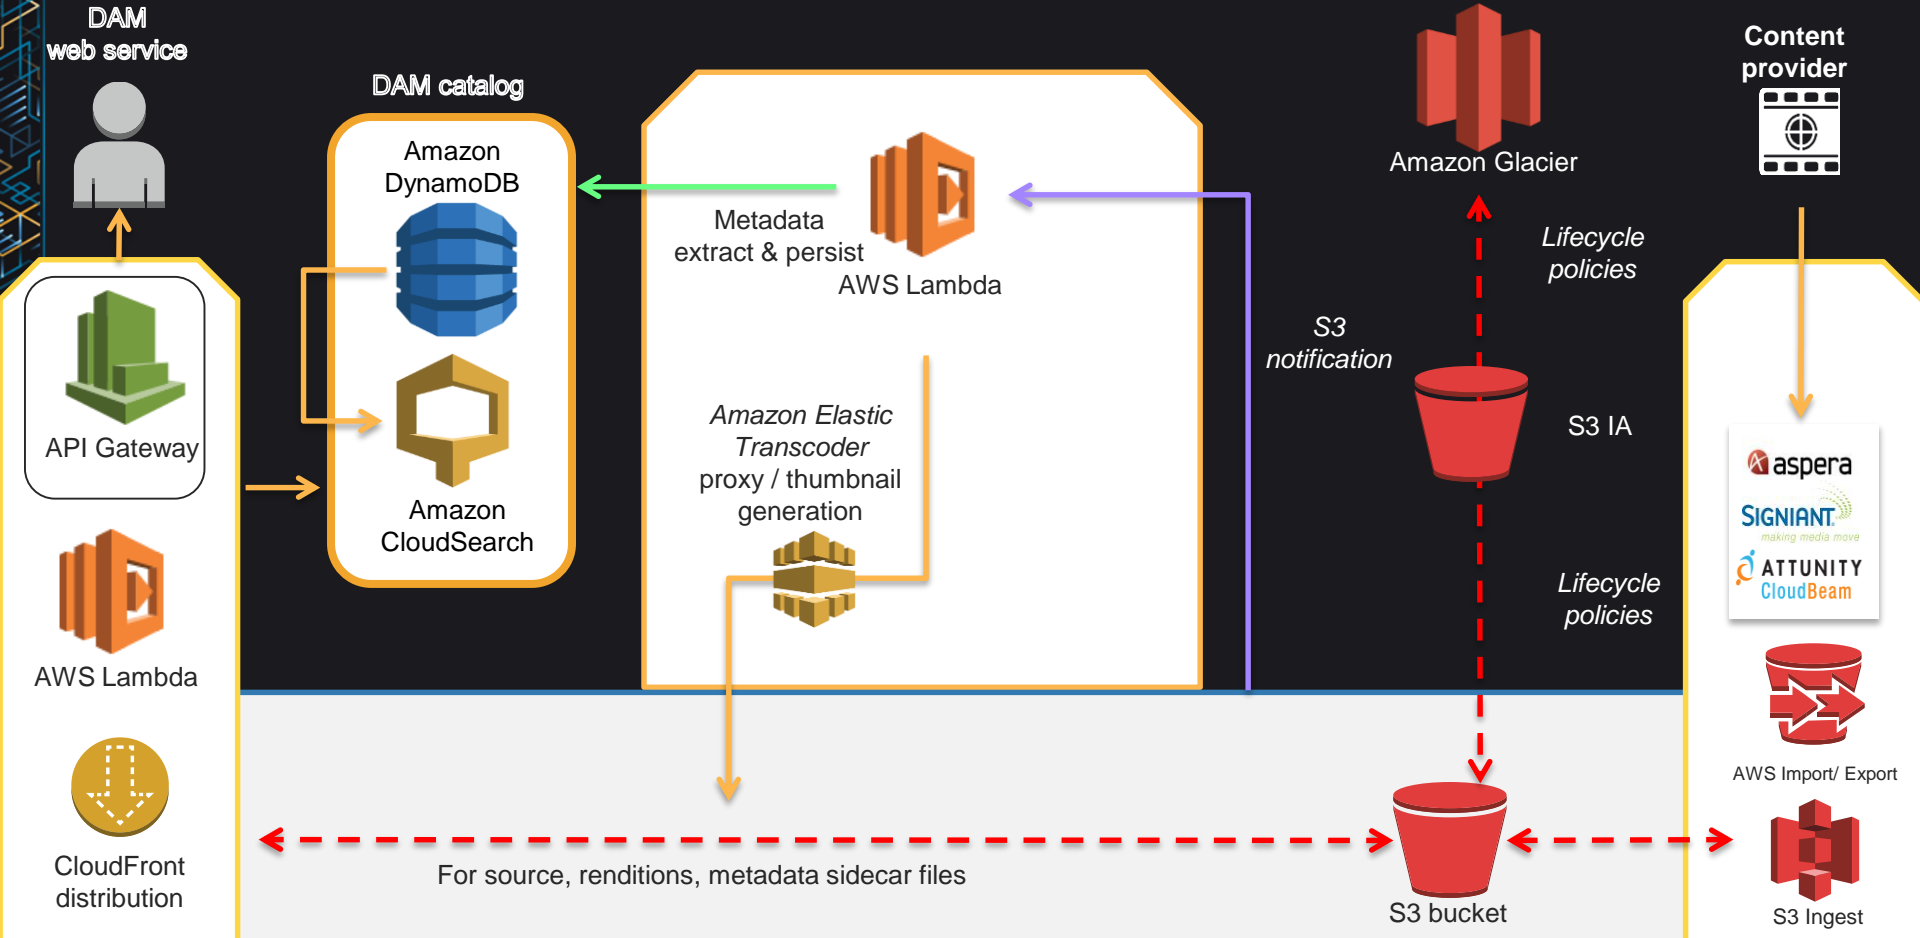
# Storage and Ingest — Serverless



# Storage and Ingest — Serverless



# Storage and Ingest — Serverless





# Encoding, Packaging, Encrypting


## Challenges

Multiple client devices

Higher quality content

Parallel, complex workflow

Uneven load distribution



# Content Processing Pipeline (Using Lambda)

Ingest



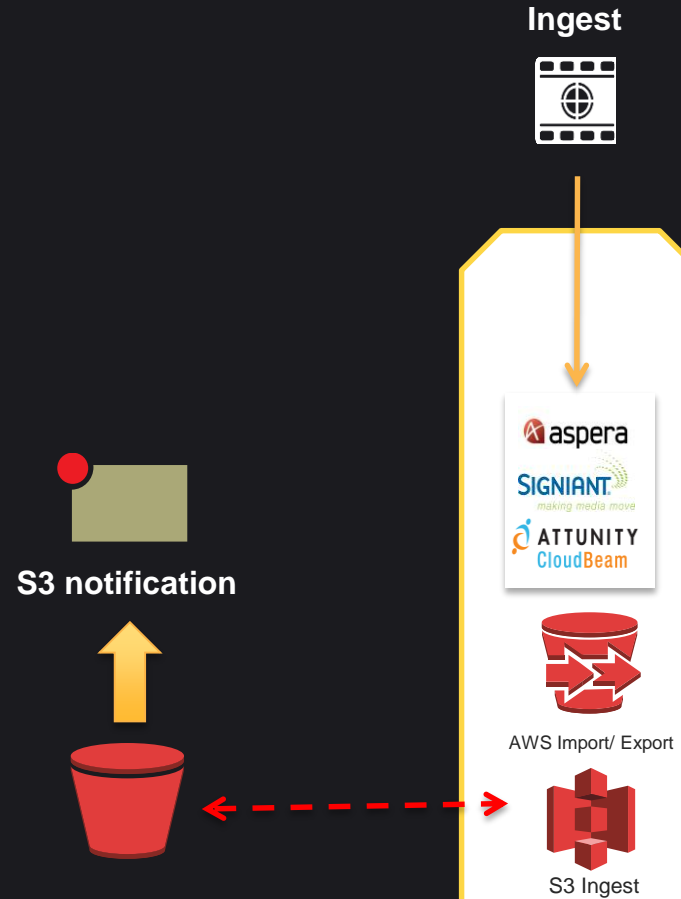
AWS Import/ Export



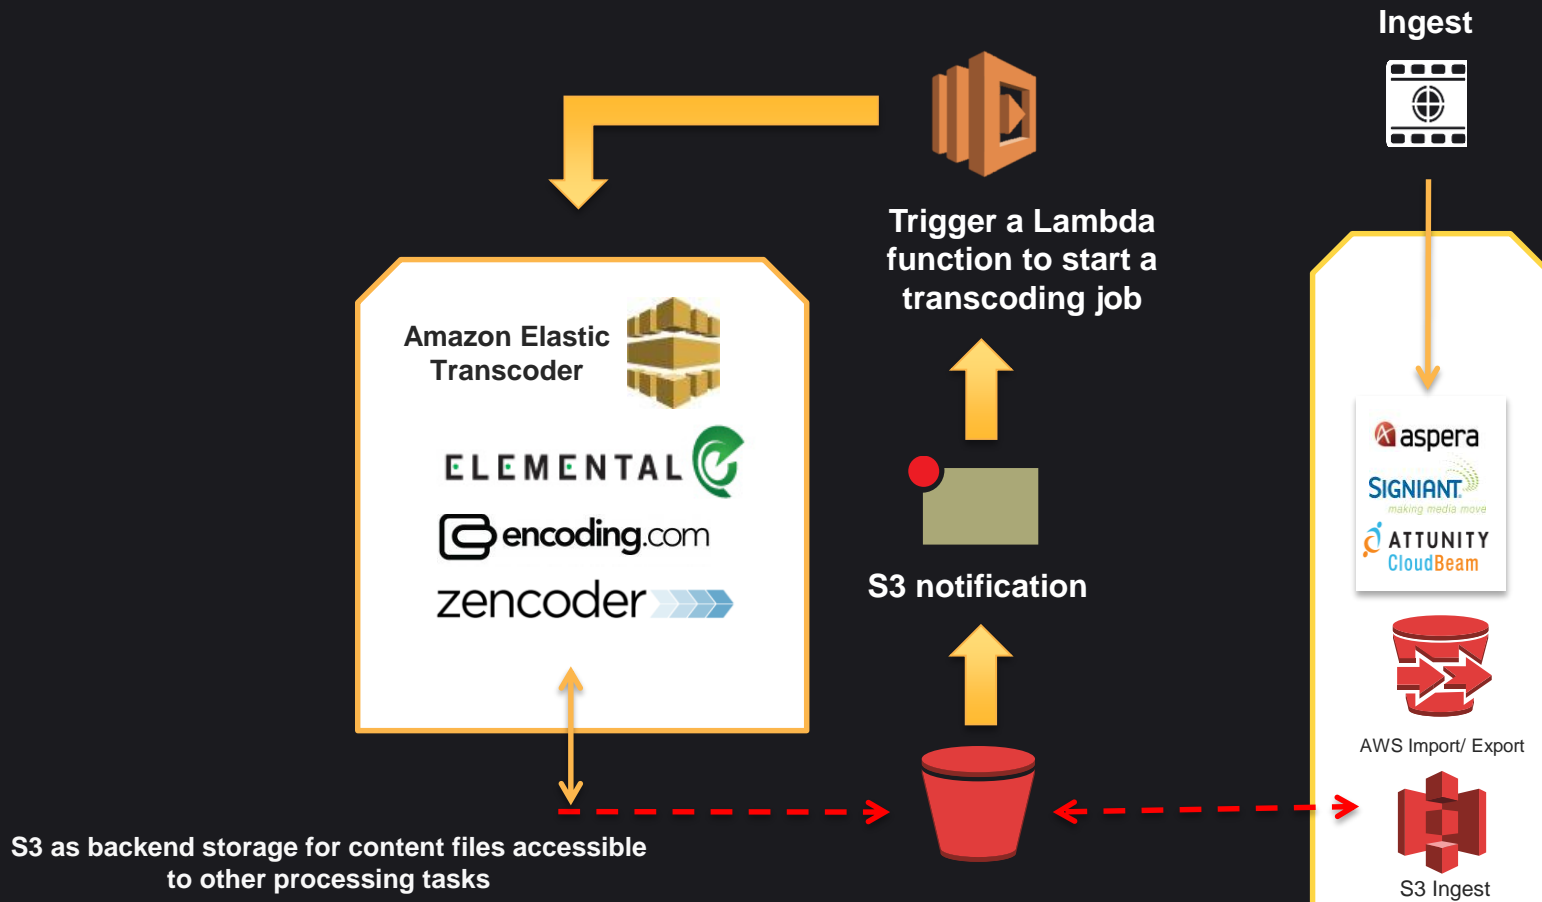
S3 Ingest



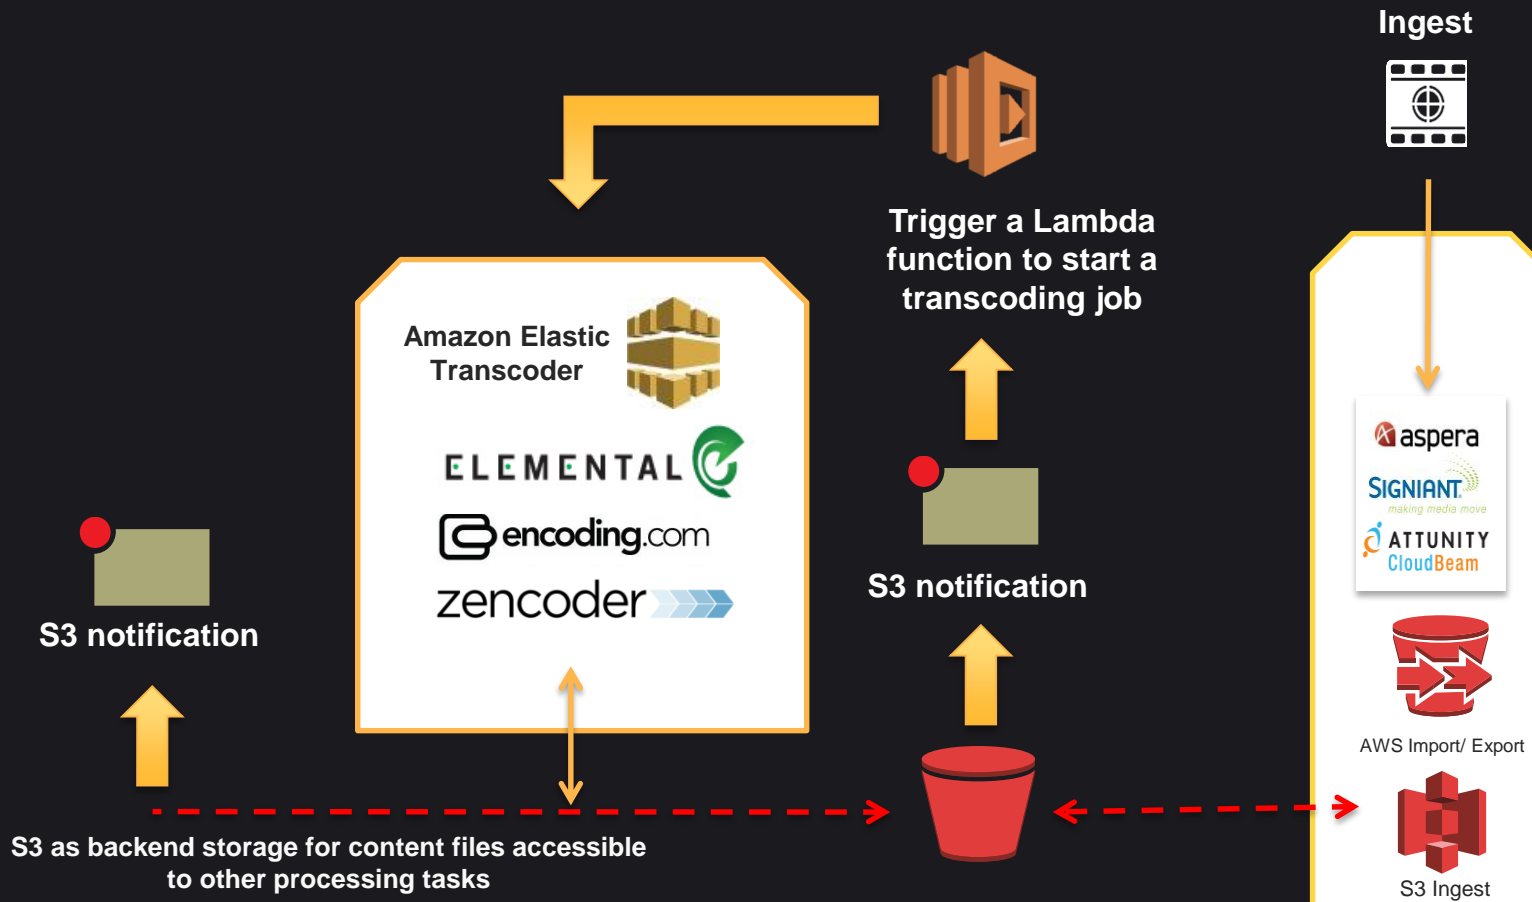
# Content Processing Pipeline (Using Lambda)



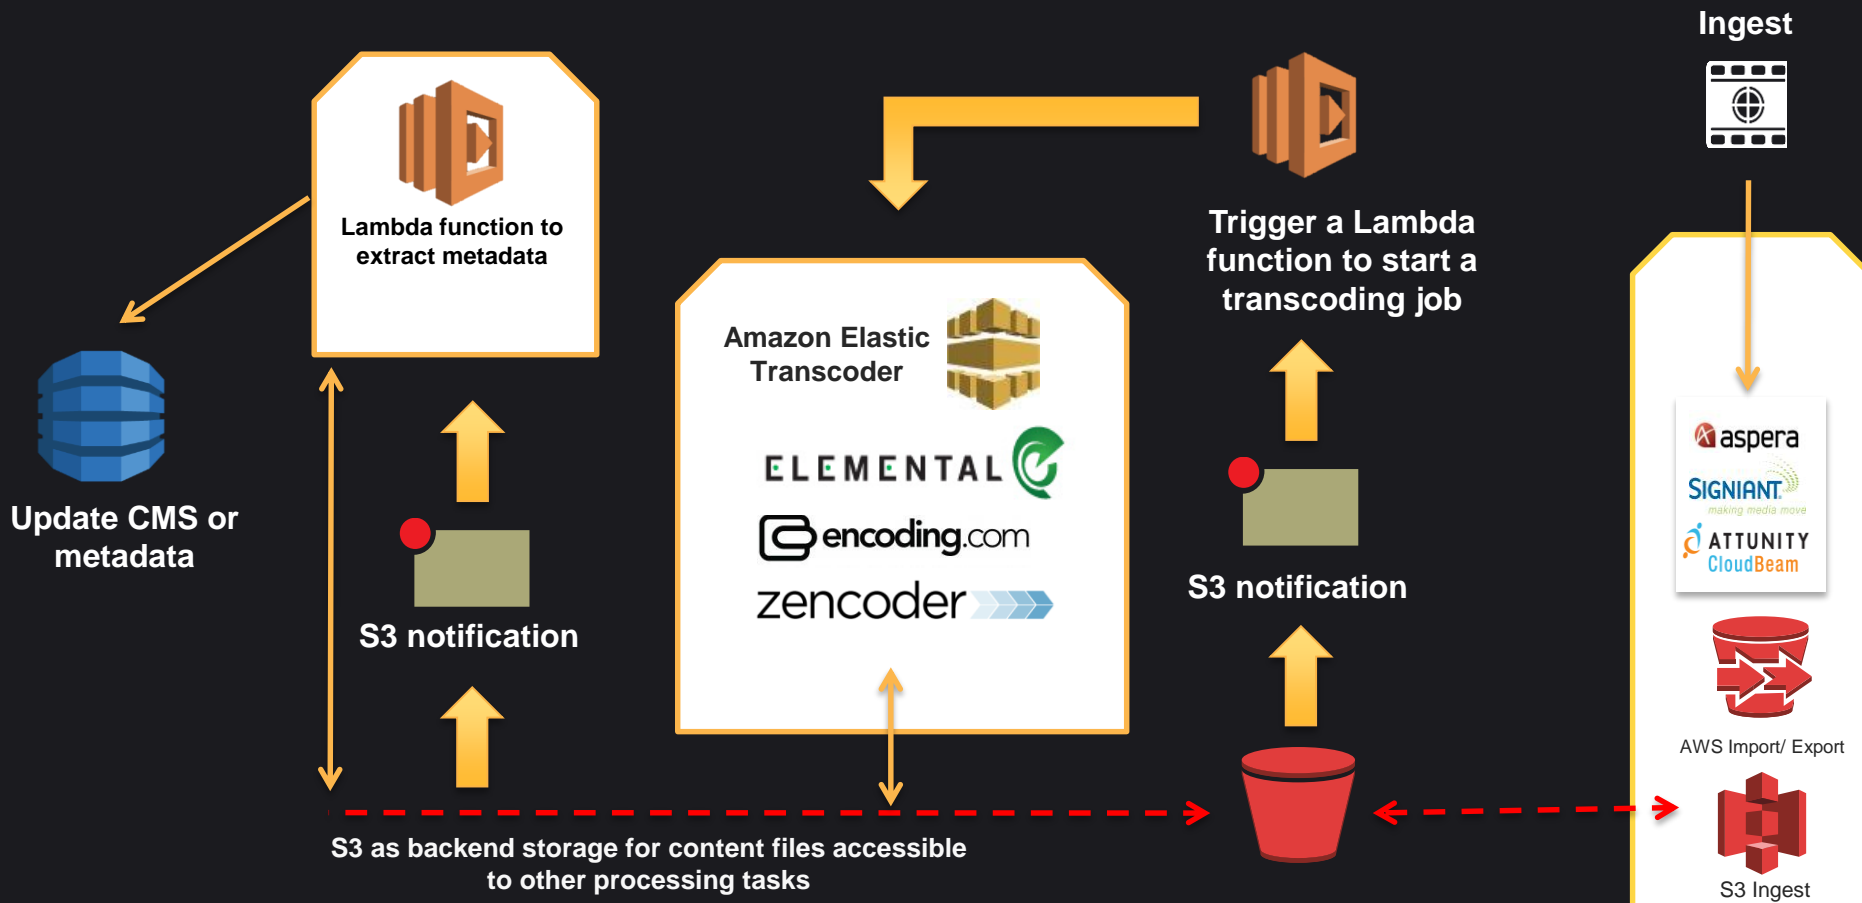
# Content Processing Pipeline (Using Lambda)



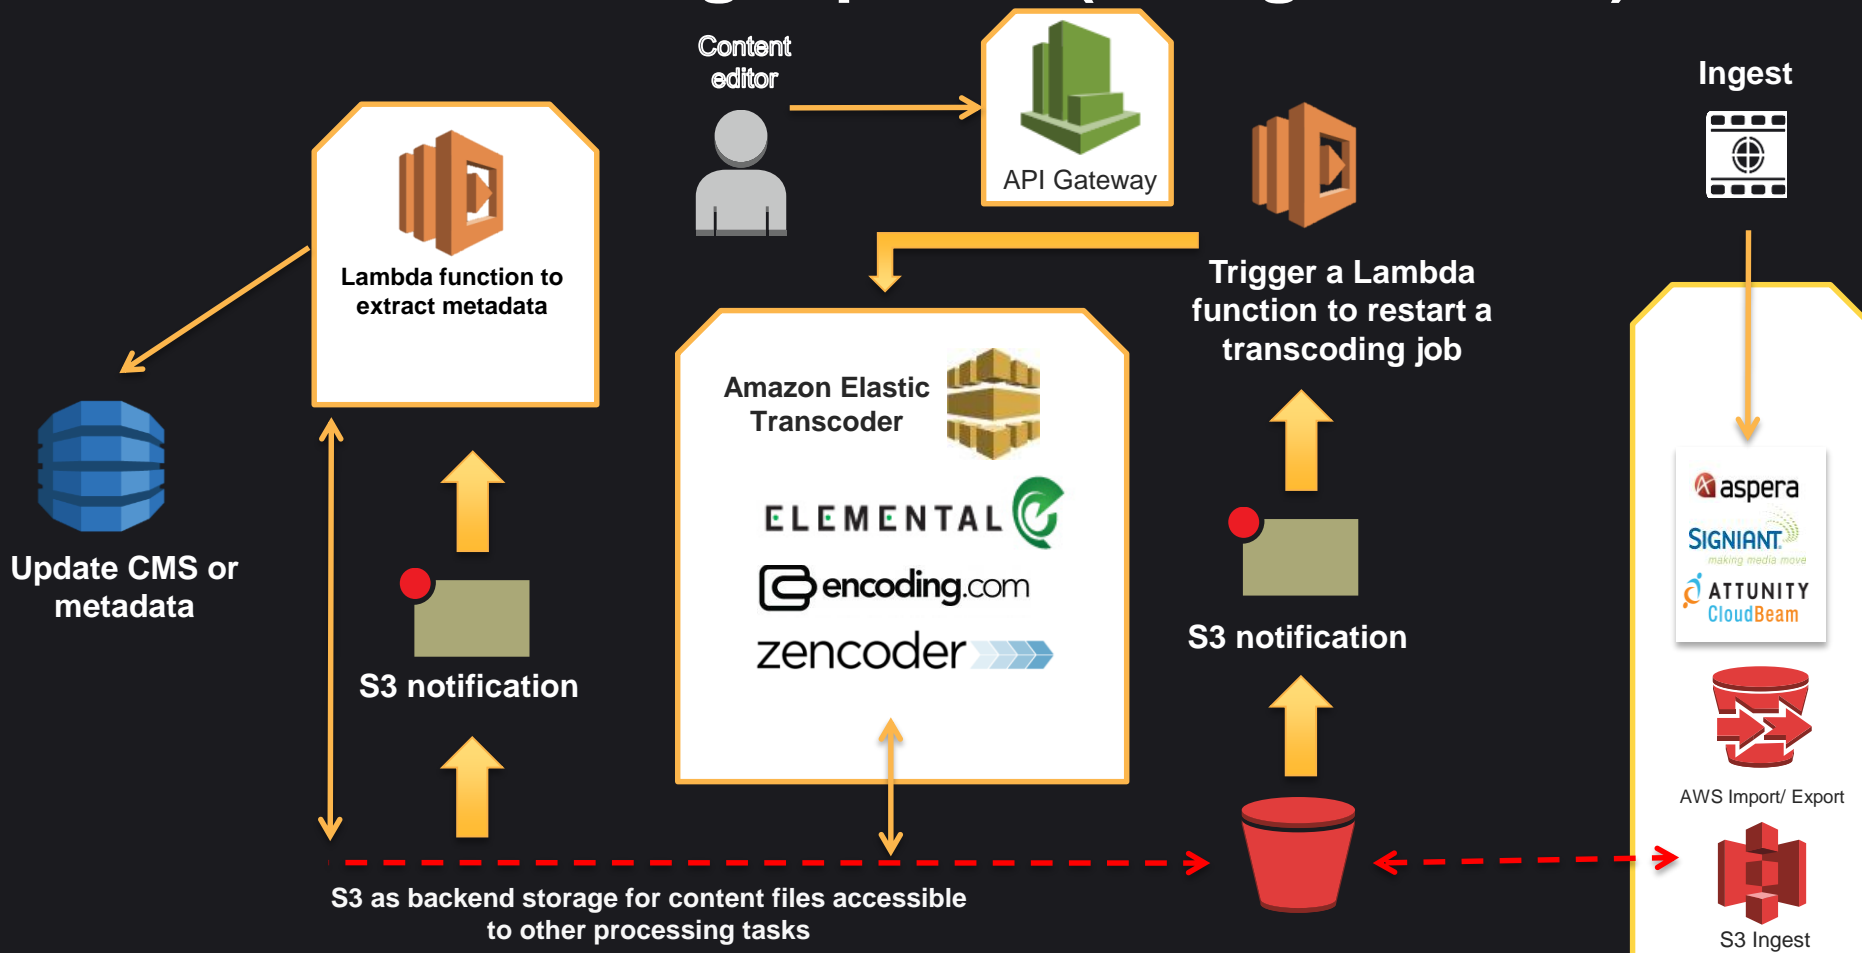
# Content Processing Pipeline (Using Lambda)



# Content Processing Pipeline (Using Lambda)



# Content Processing Pipeline (Using Lambda)




# What About Packaging?

All moving towards HTTP –  
RTMP is hard to scale

Battle of the standards –  
HLS, HDS, SmoothStream,  
MPEG-DASH

Alliance for Open Media



# And Encryption / DRM?

Still largely driven by studio requirements

Just-in-time encryption (hard)

Reusability across packaging methods (PlayReady across HLS & SmoothStream)



# Pitfalls of Content Prep

Betting big on closed standards

Technologies in vogue

Adoption is device-driven





# Avoiding Pitfalls of Content Prep

Keep your mezzanine / masters (S3 IA, Glacier)

Mix and match your encoding

Contradictory — JIT packaging and heavy caching

Remember 80/20 (95/5) rule of content




## Top Tip

CPU-based encoding

c4.8xlarge / m4.10xlarge

P-State, C-State configuration

Varies with particular encoding  
libraries – experiment and  
baseline



(Thanks PocketMath!)



pocketmath



# Content Delivery Architectures

## Live

RTMP ingest to origin

Repackaged and encrypted on the fly

DRM less common

WebRTC for 1-to-few, then RTMP/HTTP when traffic increases

## VOD / Catchup

File-based – stored on S3/CDN

Usually prepared (encrypted/packaged) before final storage

DRM by default

ISPs worried about unicast delivery

# HTTP Live Streaming (HDS & HLS)

HLS was pretty close to de facto

Space becoming disrupted again


Network considerations in the real world

# HLS and HDS

Uses “parent / child / chunk”  
model

HLS: Playlist contains  
chunklists, which contain  
chunks

HDS: Manifest contains  
bootstrap files, which contain  
fragments



# HTTP Dynamic Streaming — HDS

> GET /live/channel1.abst HTTP/1.1

< Seg1-Frag55

< Seg1-Frag56

[...]

> GET /live/channel1.abst HTTP/1.1

< Seg1-Frag1

< Seg1-Frag2

# HTTP Dynamic Streaming — HDS

Use fragment name alignment in the event of a republish

Ensure that fragments are aligned **across bitrates**

Players will error otherwise






## Top Tip

HDS bootstraps are binary

To view, you must decode  
them first

<https://bitly.com/abstdecoder>



(Thanks SwiftServe!)

# HTTP Live Streaming — HLS

> GET /live/channel1.m3u8 HTTP/1.1

< HTTP/1.1 200 OK

< Date: Tue, 15 Sep 2015 13:37:56 GMT

< Server: Apache

< Last-Modified: Mon, 03 Aug 2015 3:14:15 GMT

< Accept-Ranges: bytes

< Content-Length: 219

# HTTP Live Streaming — HLS

> GET /live/channel1.m3u8 HTTP/1.1

< HTTP/1.1 200 OK

< Date: Tue, 15 Sep 2015 13:37:56 GMT

< Server: Apache

< Last-Modified: Mon, 03 Aug 2015 3:14:15 GMT

< Accept-Ranges: bytes

< Content-Length: 219

< Cache-Control: max-age=5

# HTTP Live Streaming

Add a no, or short age, cache header

Use segment name randomisation

Use Player metrics to detect problem ISPs / CDNs




# Multi-CDN

# CDN Selector

Build vs. buy

Excellent offerings on the market

Metadata often tightly coupled with platform, so tricky to use a third party



# CDN Selection Methods

## DNS-based

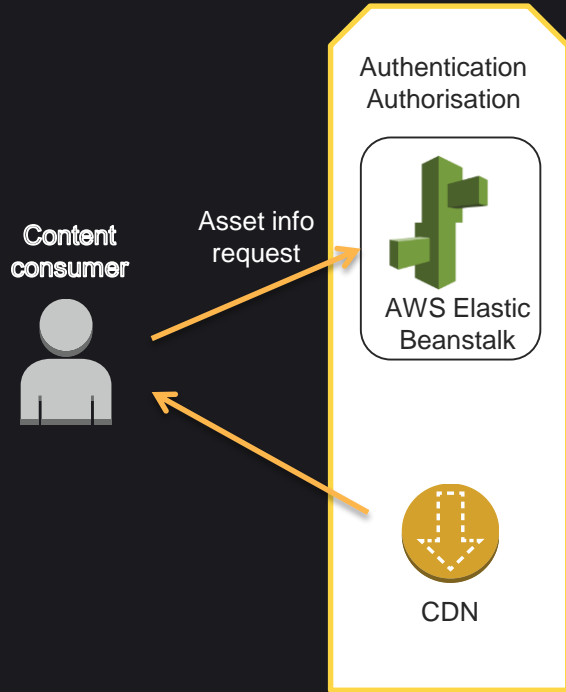
Geo / latency / intelligent routing at DNS level

## Asset sharding

50% of assets on CDN A, 50% on CDN B

## CDN-aware asset info service

# CDN-aware Asset Info Service



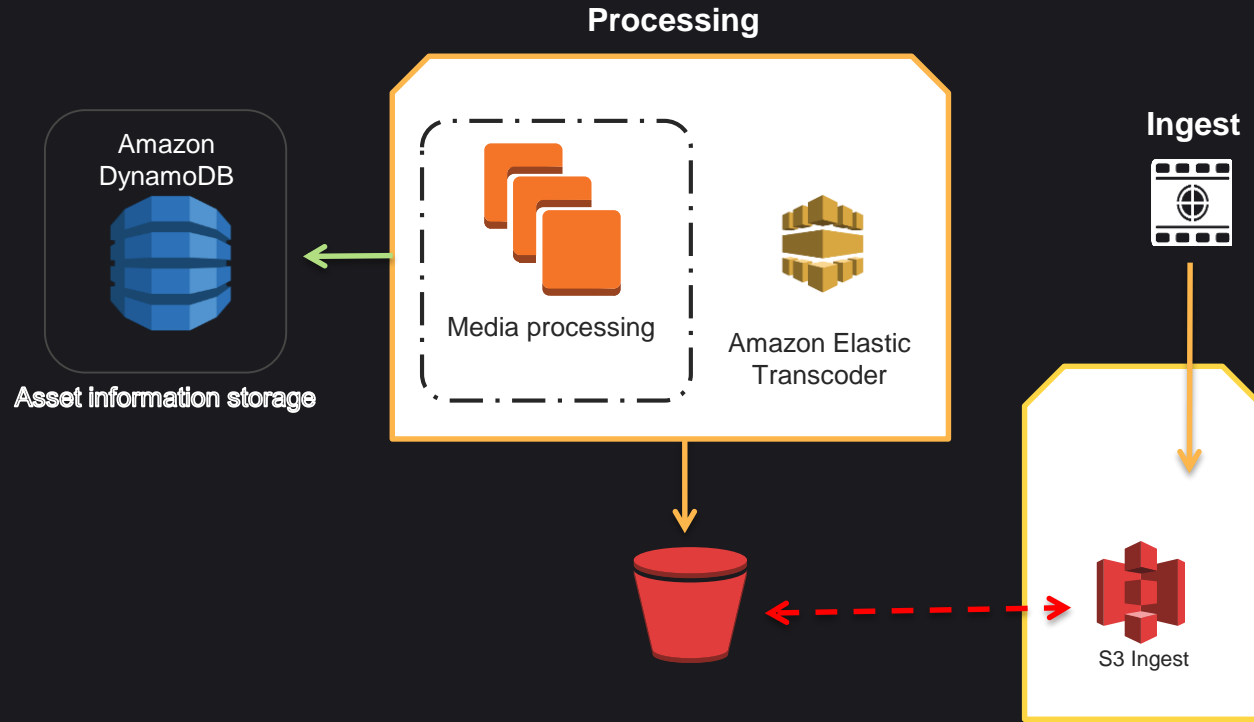


# CDN-aware Asset Info Service

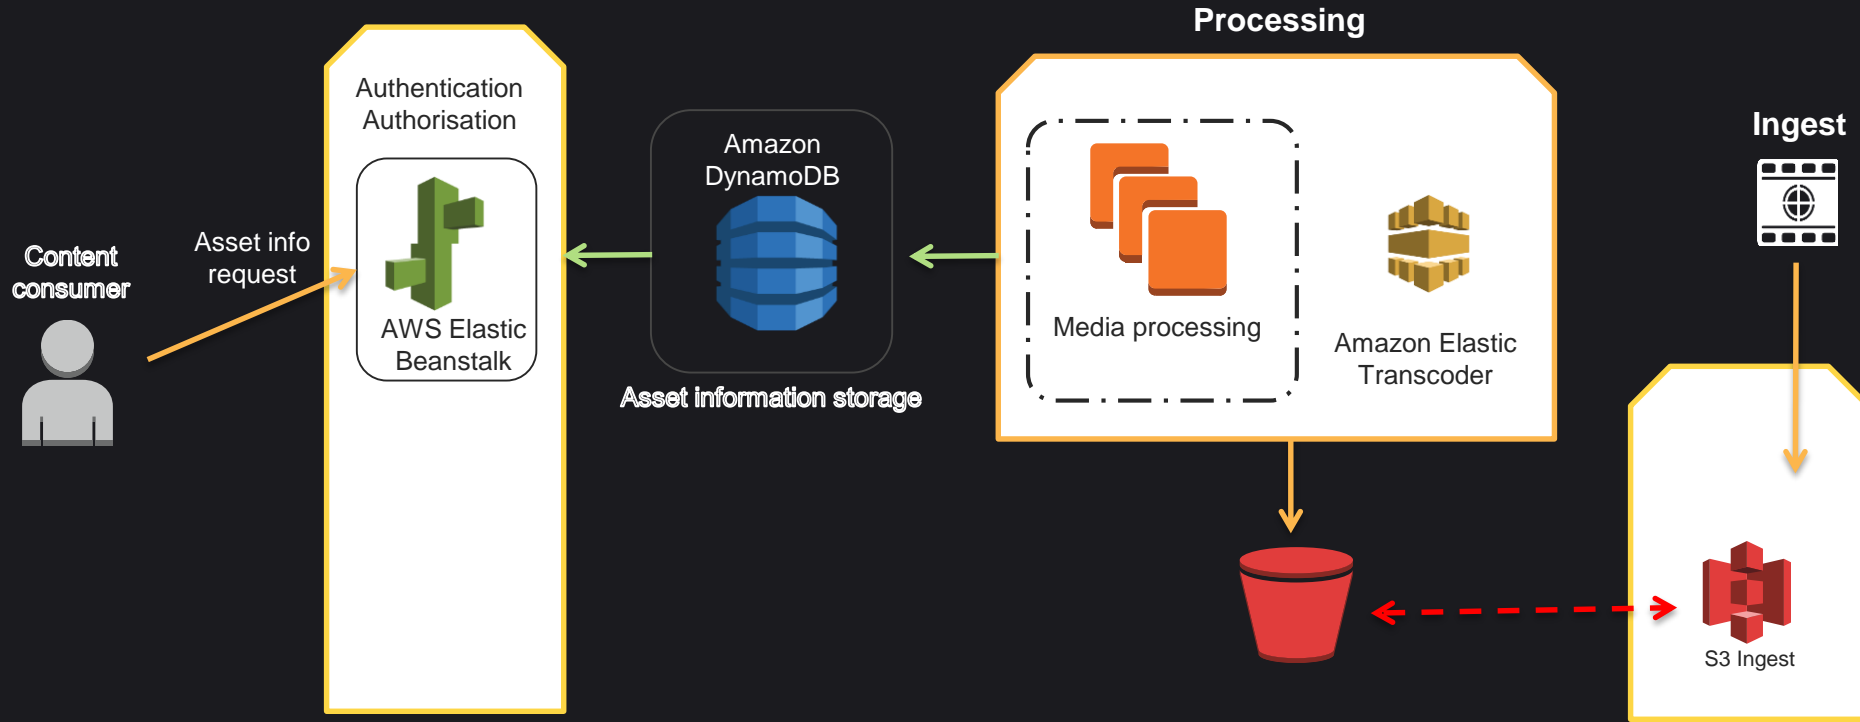
> GET /asset/31337 HTTP/1.1

```
< assetUrl: "http://cdn-a.alexjs.im/vod/31337.m3u8"  
< adProvider: "alexjsAds"  
< countryCode: "[im, sg, id]"
```

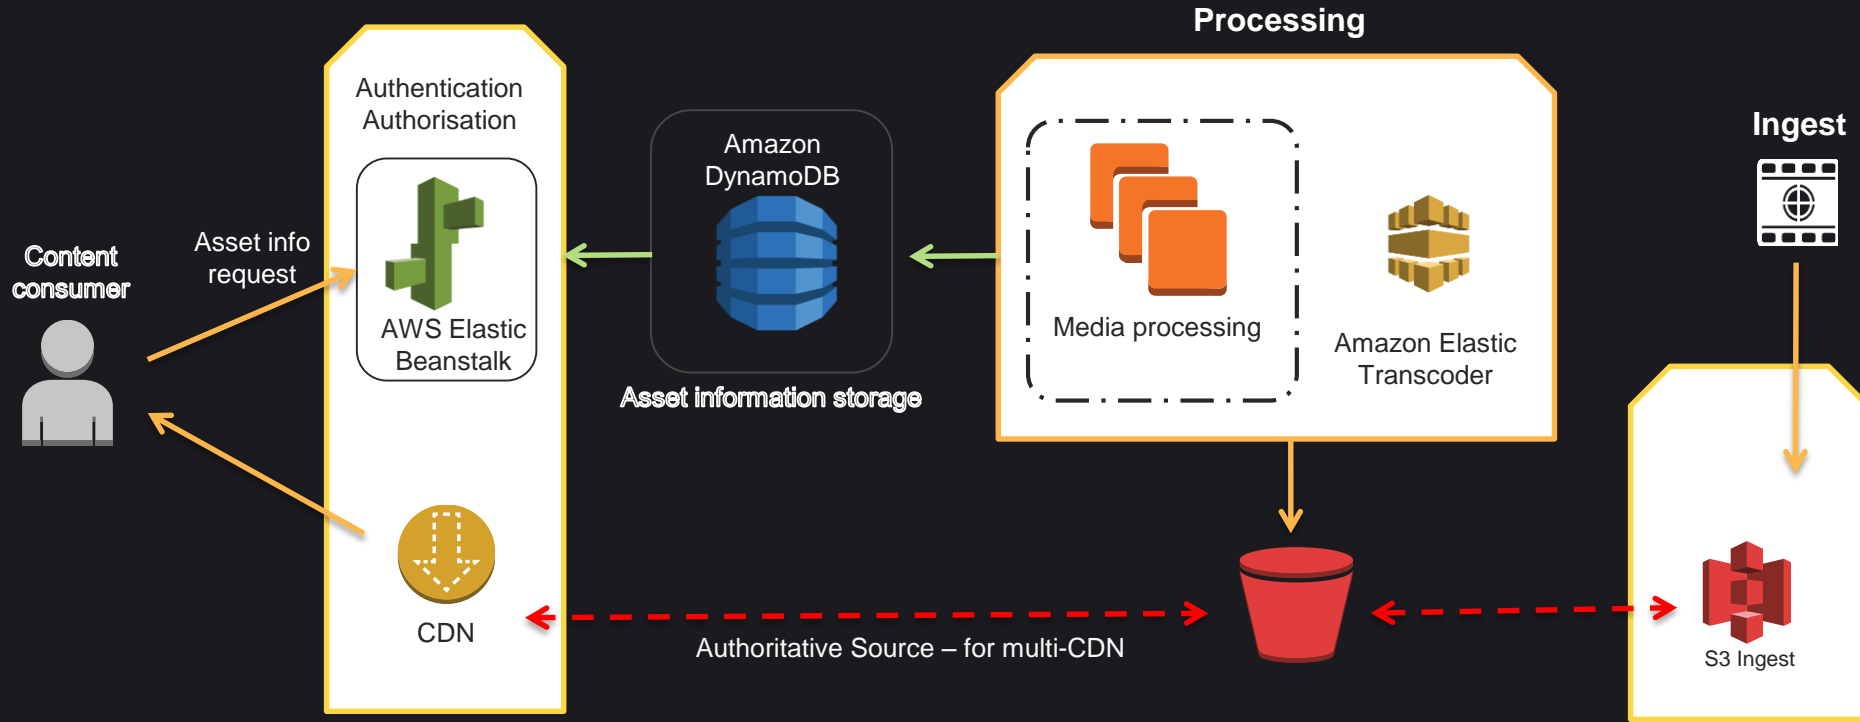
# CDN-aware Asset Info Service



# CDN-aware Asset Info Service

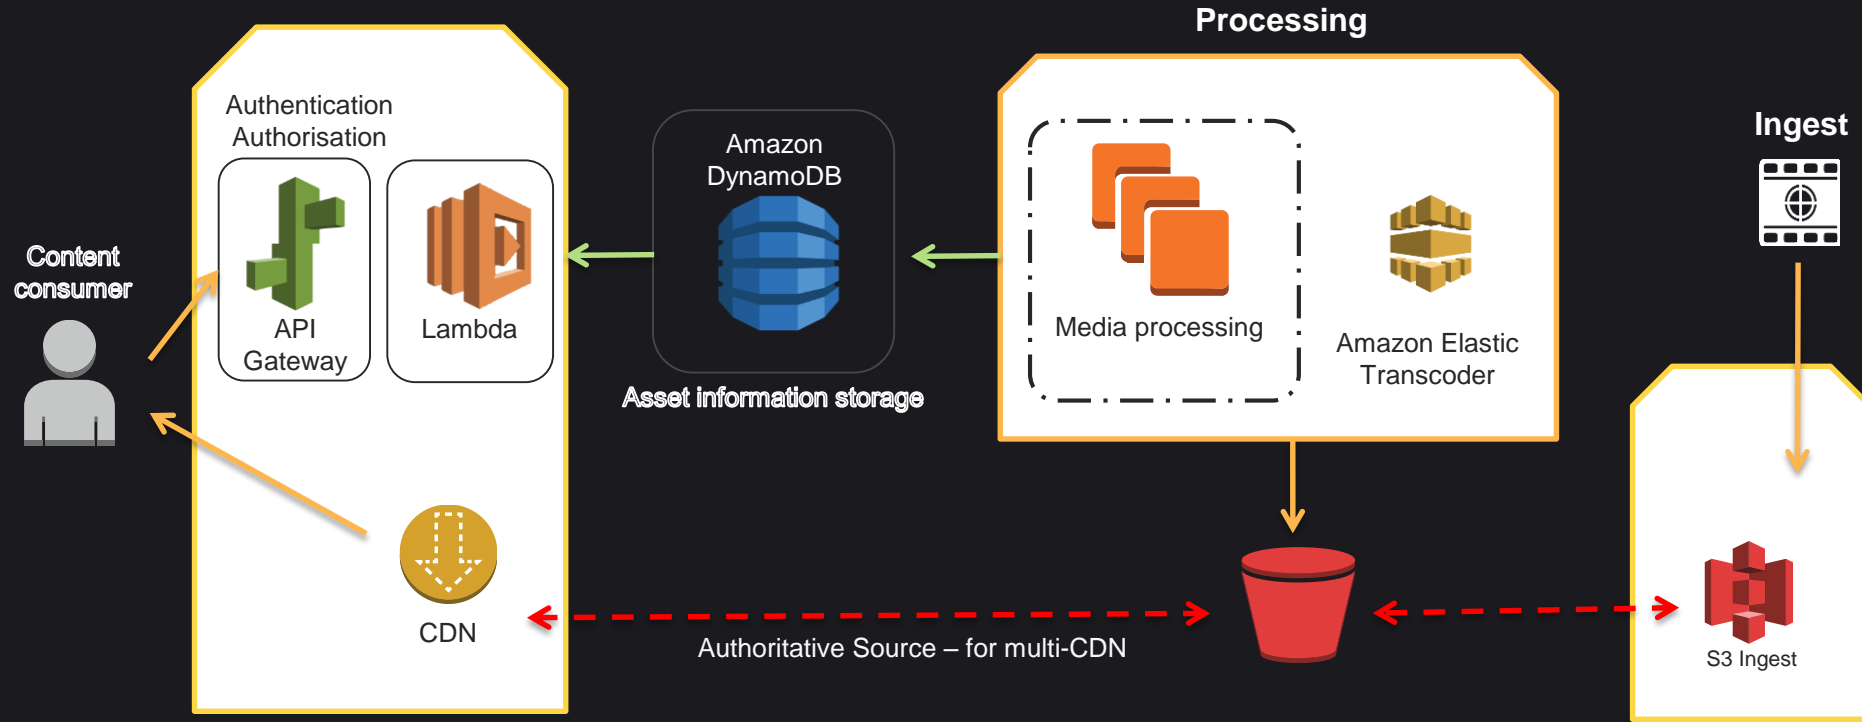


# CDN-aware Asset Info Service



# CDN-aware Asset Info Service

Perfect for microservices – and for serverless computing







# Concurrency Management (and More)

# Concurrency

## Problem

Studio mandate on stream concurrency

Cross-account sharing


## Approach

Player heartbeat

Subscription-based thresholds



# Subscriber Concurrency Management



Heartbeat sent from every player at regular intervals

Supports both native and web-based players

# Subscriber Concurrency Management

```
{  
  assetId: "d6f9fe" // Programme ID  
  userId: "33114220875dc" // Token / User  
ID timeStamp: "T00:00:05" // Progress  
  deviceId: "93d2d4fef95cb" // Fingerprint  
  deviceType: "Amazon Fire" // From API  
}
```

(Pseudo JSON)

Client Player  
& Mobile  
Apps

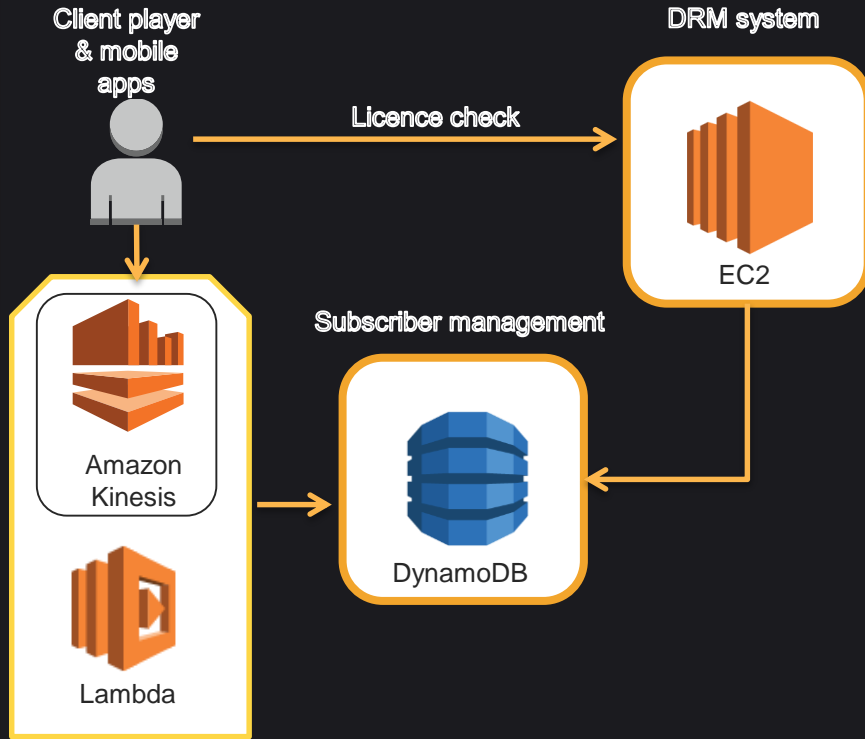


Amazon  
Kinesis



Lambda

# Subscriber Concurrency Management



Heartbeat received by Amazon Kinesis

Data fed into DynamoDB

DRM licence server reads DynamoDB table

Decision made

# Subscriber Concurrency Management

Client player  
& mobile  
apps



Subscriber management



Data warehouse



Track drop-off

Highlight popular content

Feed back into content  
development

# Subscriber Concurrency Management

Client Player  
& Mobile  
Apps



Subscriber management




Data warehouse



Syndicate data to ads server

Recommendations-based  
demographic

# Subscriber Concurrency Management



Cross-device

Play / pause / resume

No servers needed

Client-side “stop” can be enabled, too (non-DRM based)


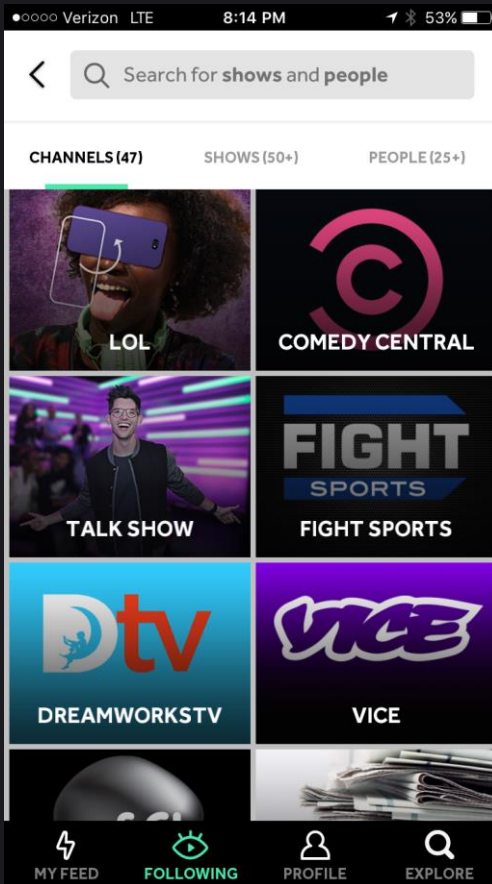



**Go90 — Verizon**

# Go90



# go90.com, #go90









# Go90 — from 30,000 ft

100+ micro services/systems

10+ deployments per week

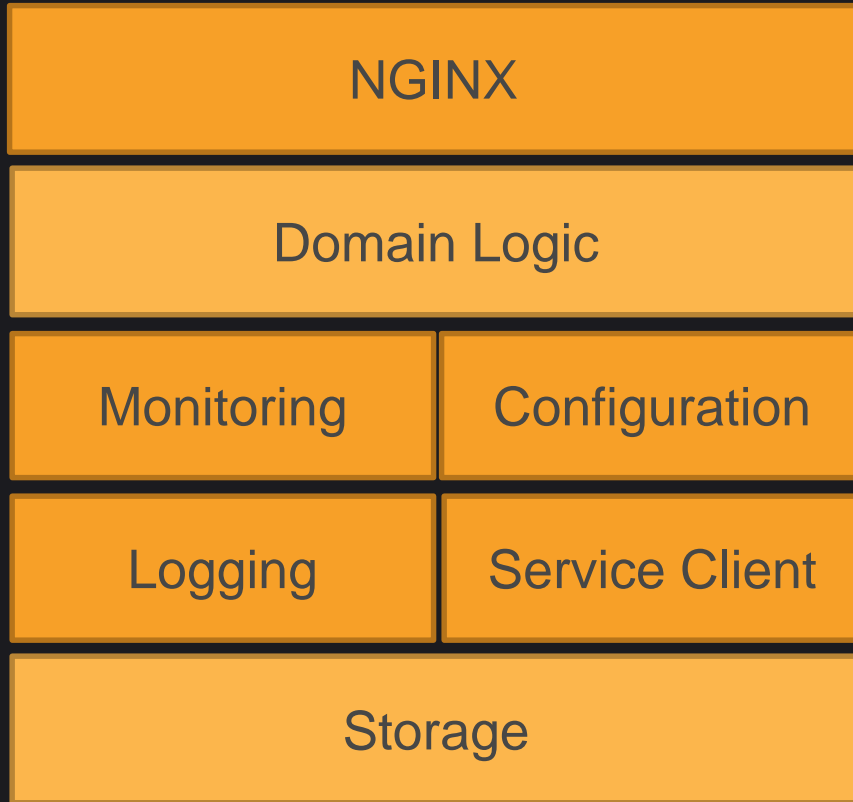
Multiple clients, large target user base

High reliability/availability, low latency, superior user experience



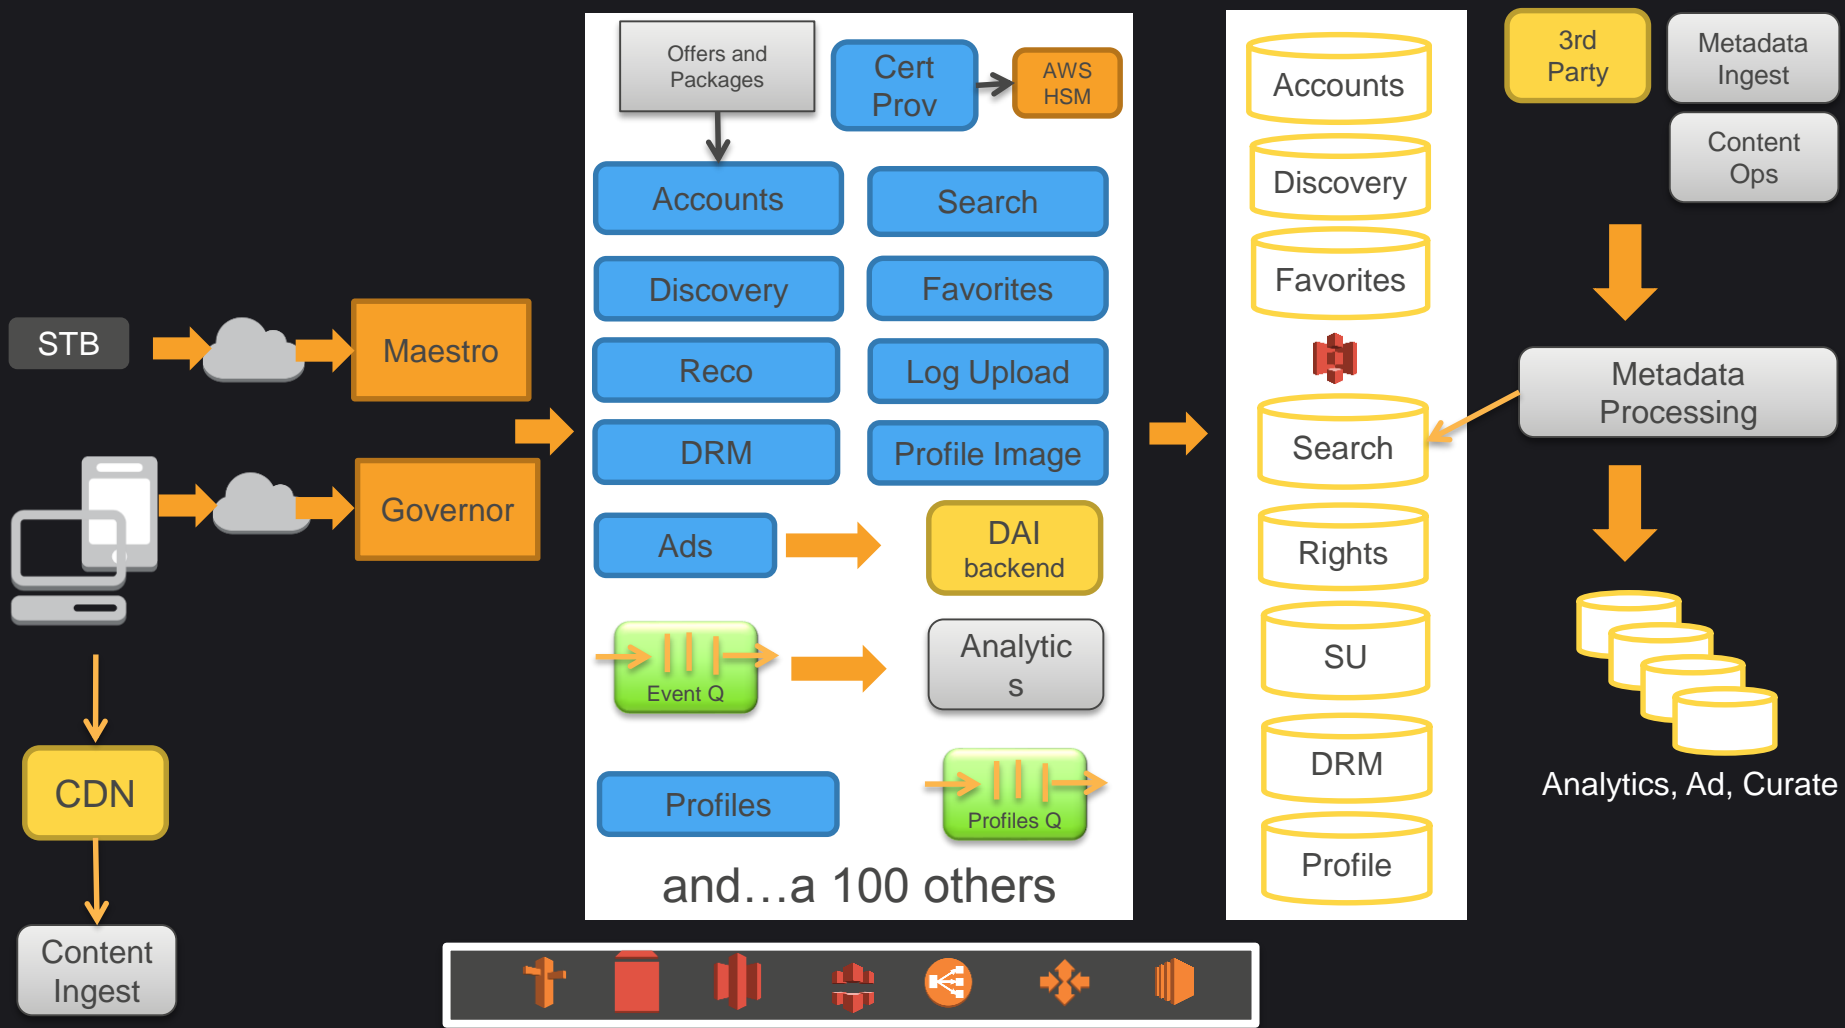
# APIs — Edge Services

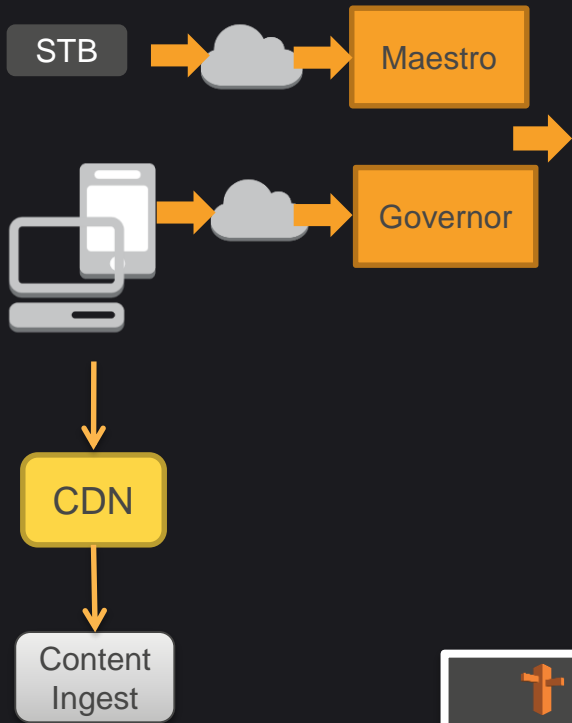
# Typical Service Stack

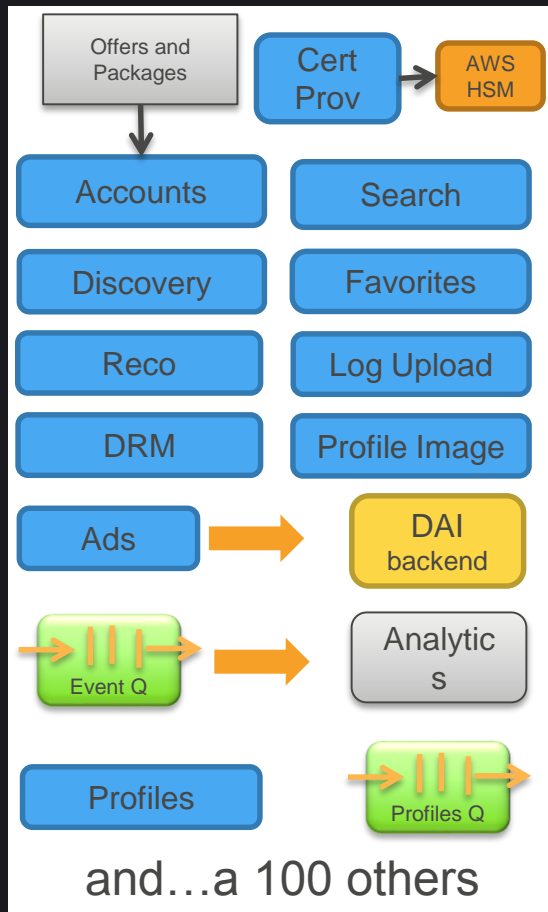



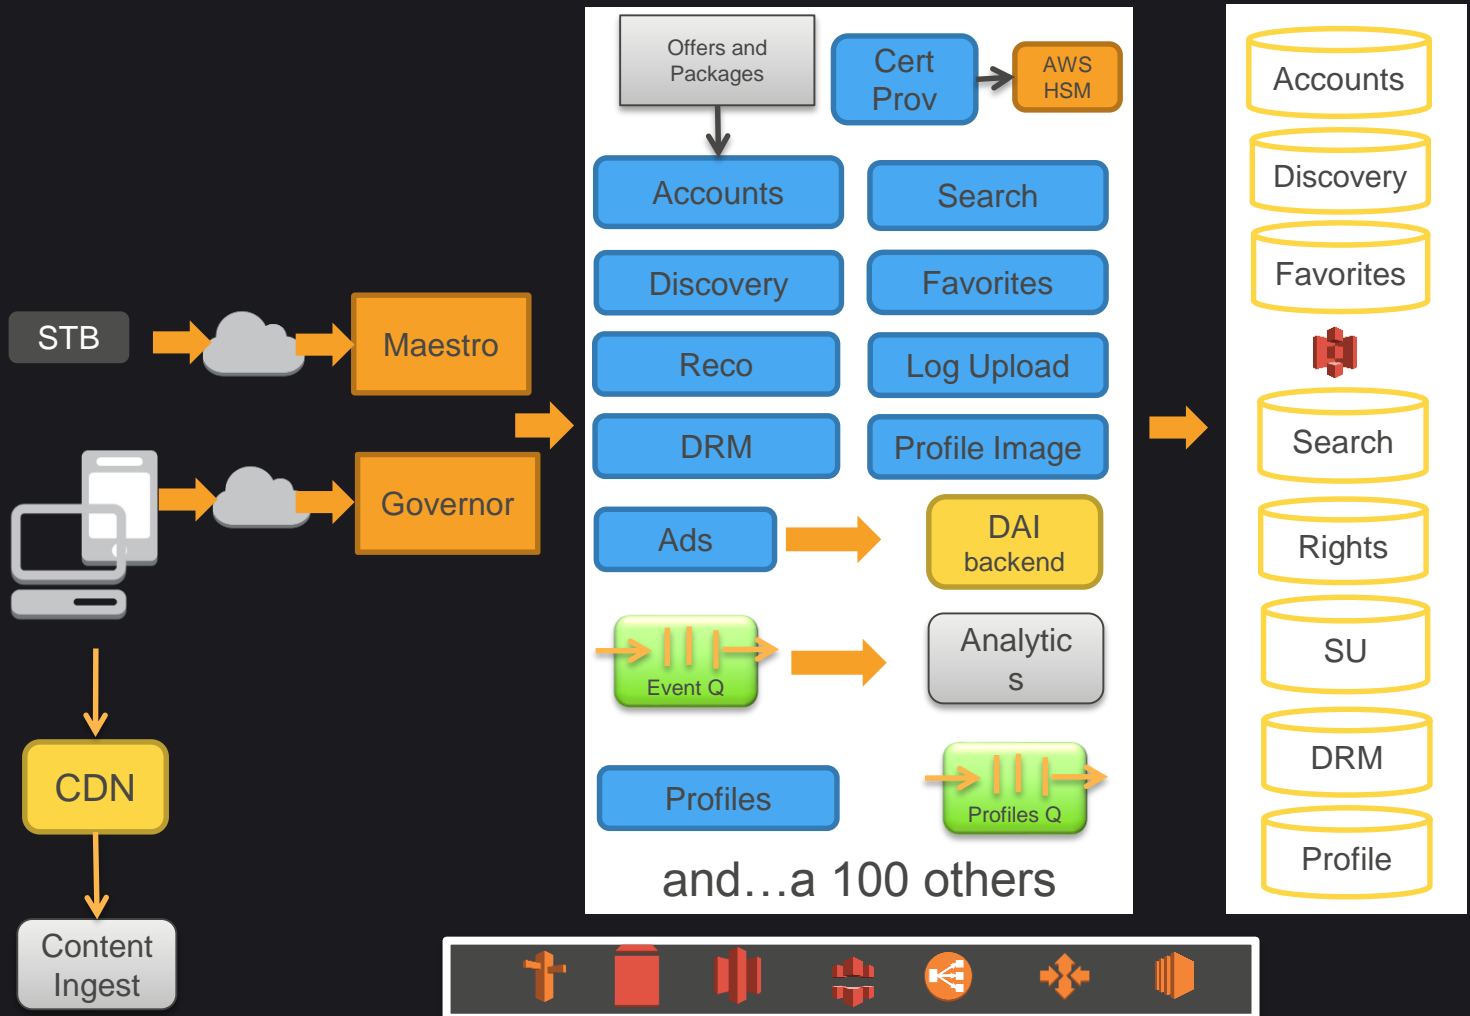
**verizon**<sup>v</sup>

<https://github.com/oncue>

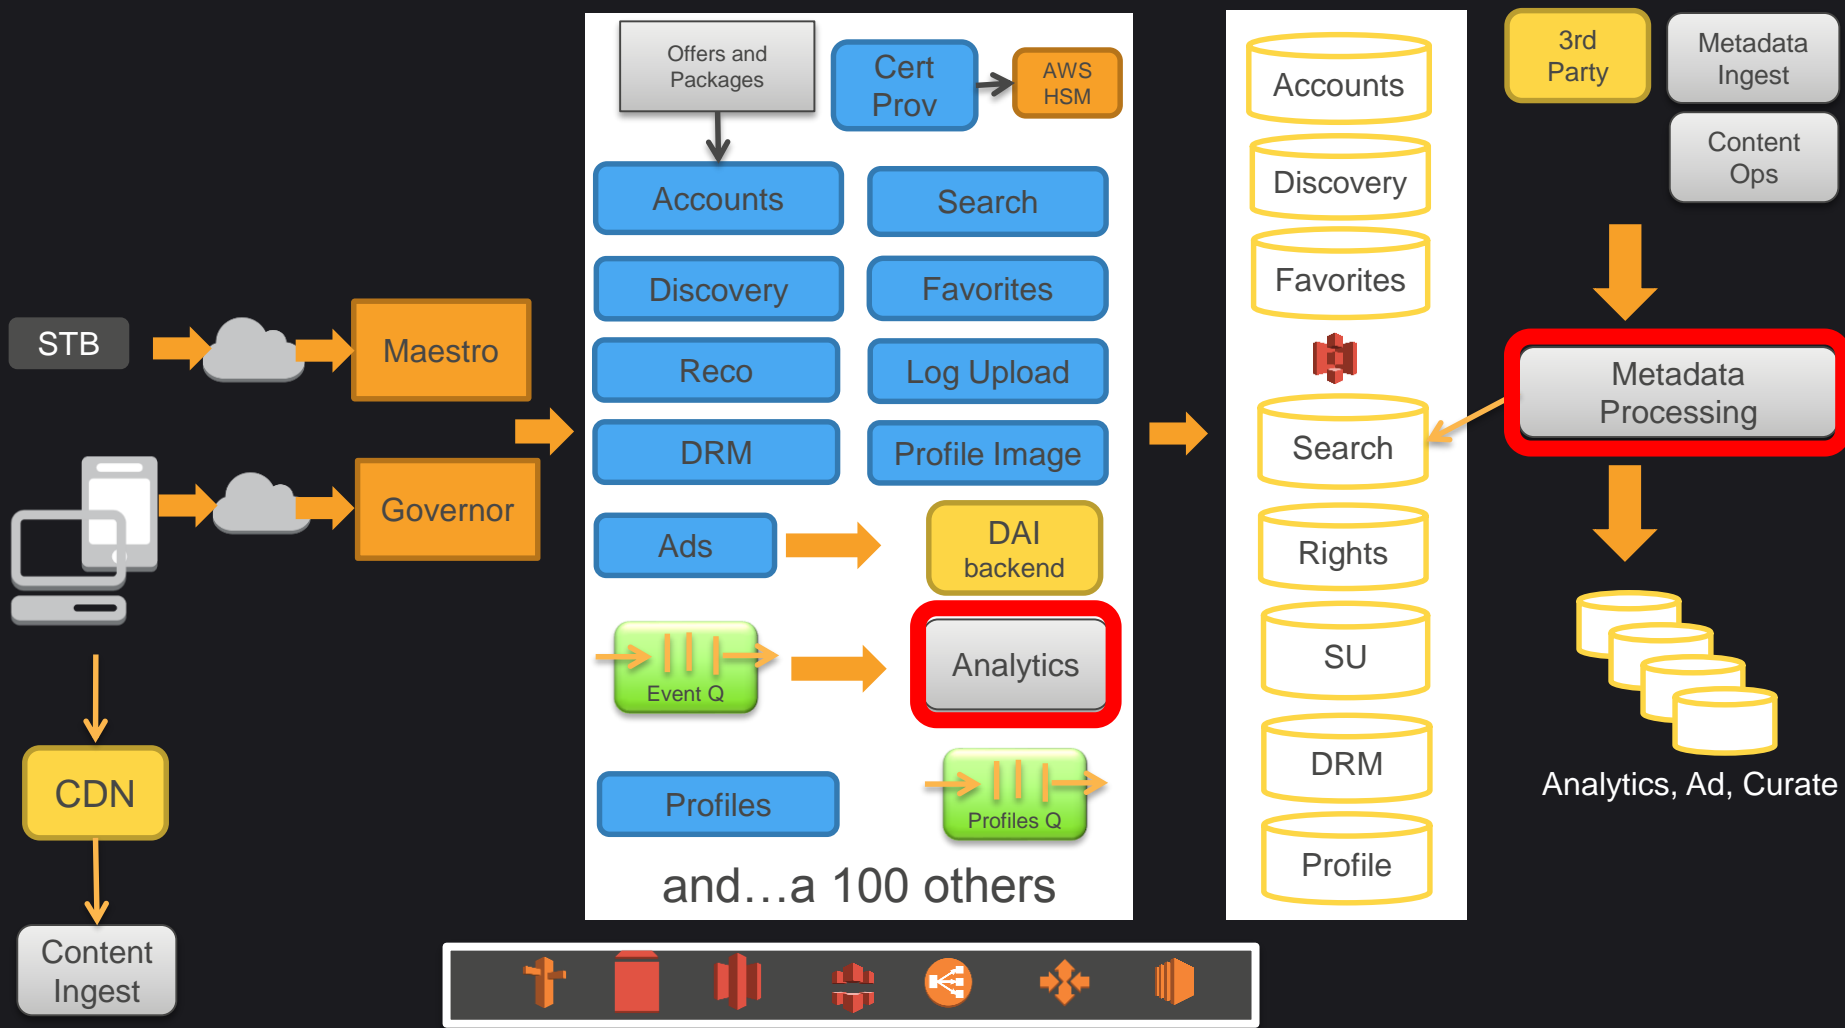














# Metadata — Ingest and Storage



Contact the Filmmakers on IMDbPro »

## The Martian (2015)



PG-13 | 141 min | Action, Adventure, Sci-Fi | 2 October 2015 (USA)



**Your rating:** ★★★★★★★★ -/10

Ratings: **8.6**/10 from 1,735 users Metascore: 76/100

Reviews: 4 user | 32 critic | 12 from Metacritic.com

During a manned mission to Mars, Astronaut Mark Watney is presumed dead after a fierce storm and left behind by his crew. But Watney has survived and finds himself stranded and alone on the hostile planet. With only meager supplies, he must draw upon his ingenuity, wit and spirit to subsist and find a way to signal to Earth that he is alive.

**Director:** Ridley Scott

**Writers:** Drew Goddard (screenplay), Andy Weir (book)

**Stars:** Matt Damon, Jessica Chastain, Kristen Wiig |

[See full cast and crew »](#)

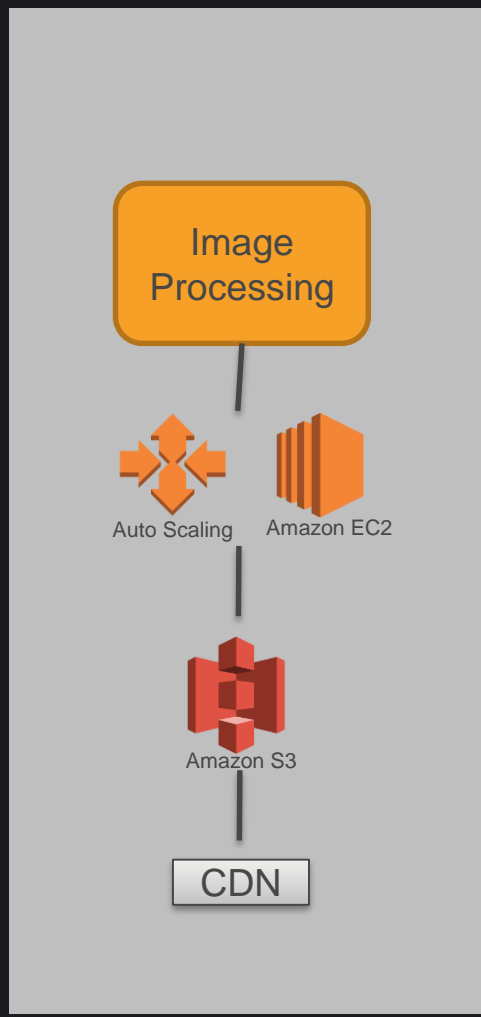
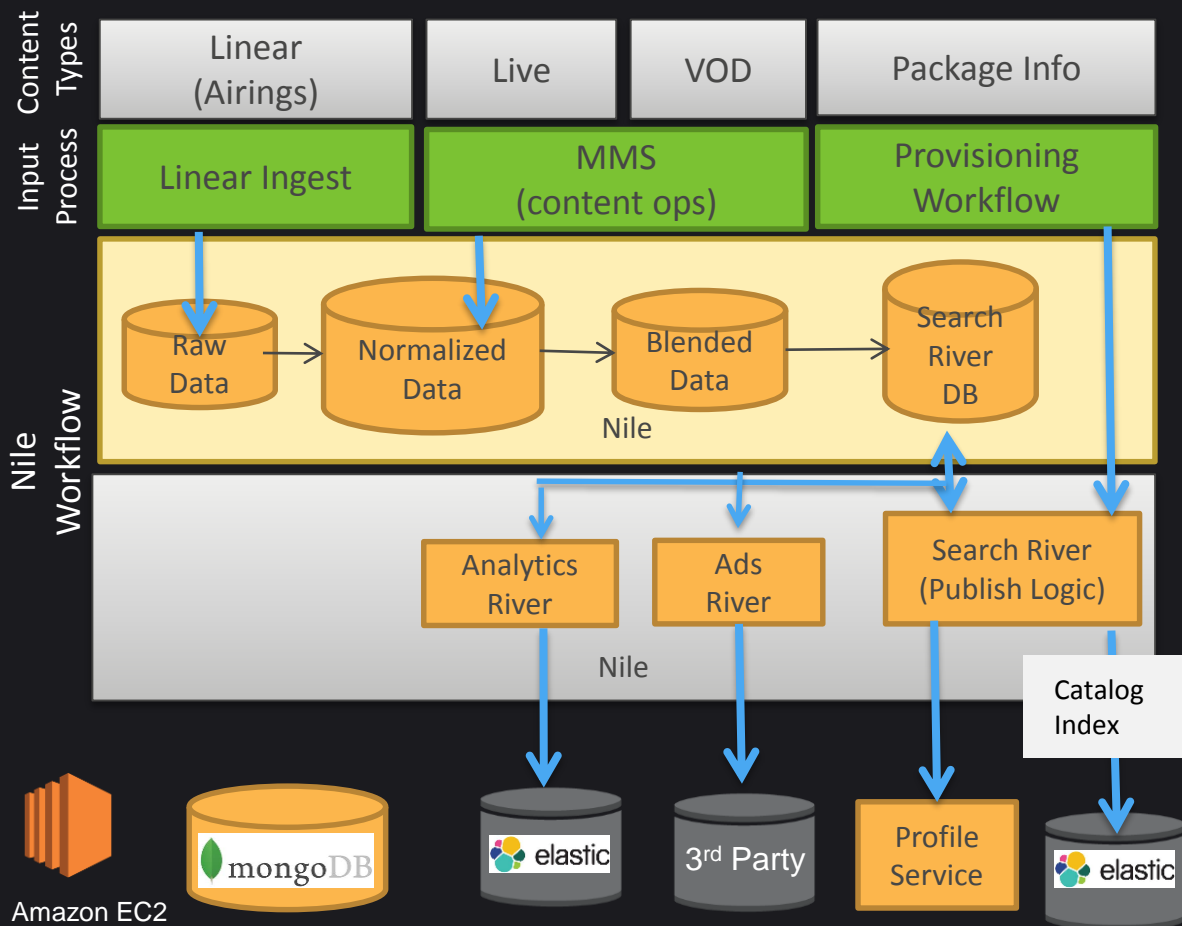
+ Watchlist



Watch Trailer

Share...

Courtesy: IMDB



Content  
Types

Linear  
(Airings)

Live

VOD

Package Info



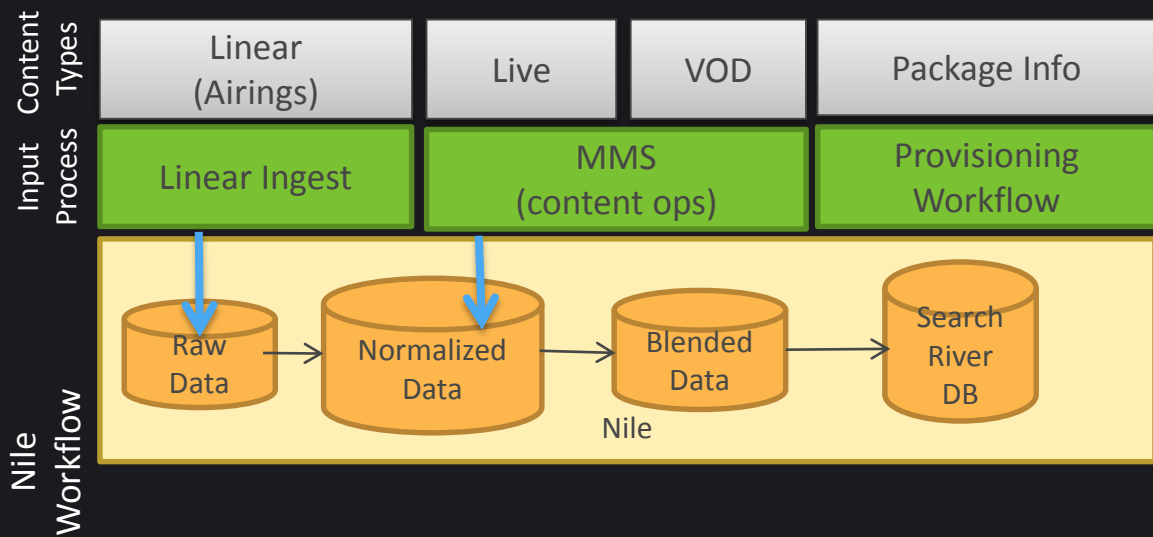
Amazon EC2

Input Content  
Process Types

Linear (Airings)	Live	VOD	Package Info
Linear Ingest	MMS (content ops)		Provisioning Workflow

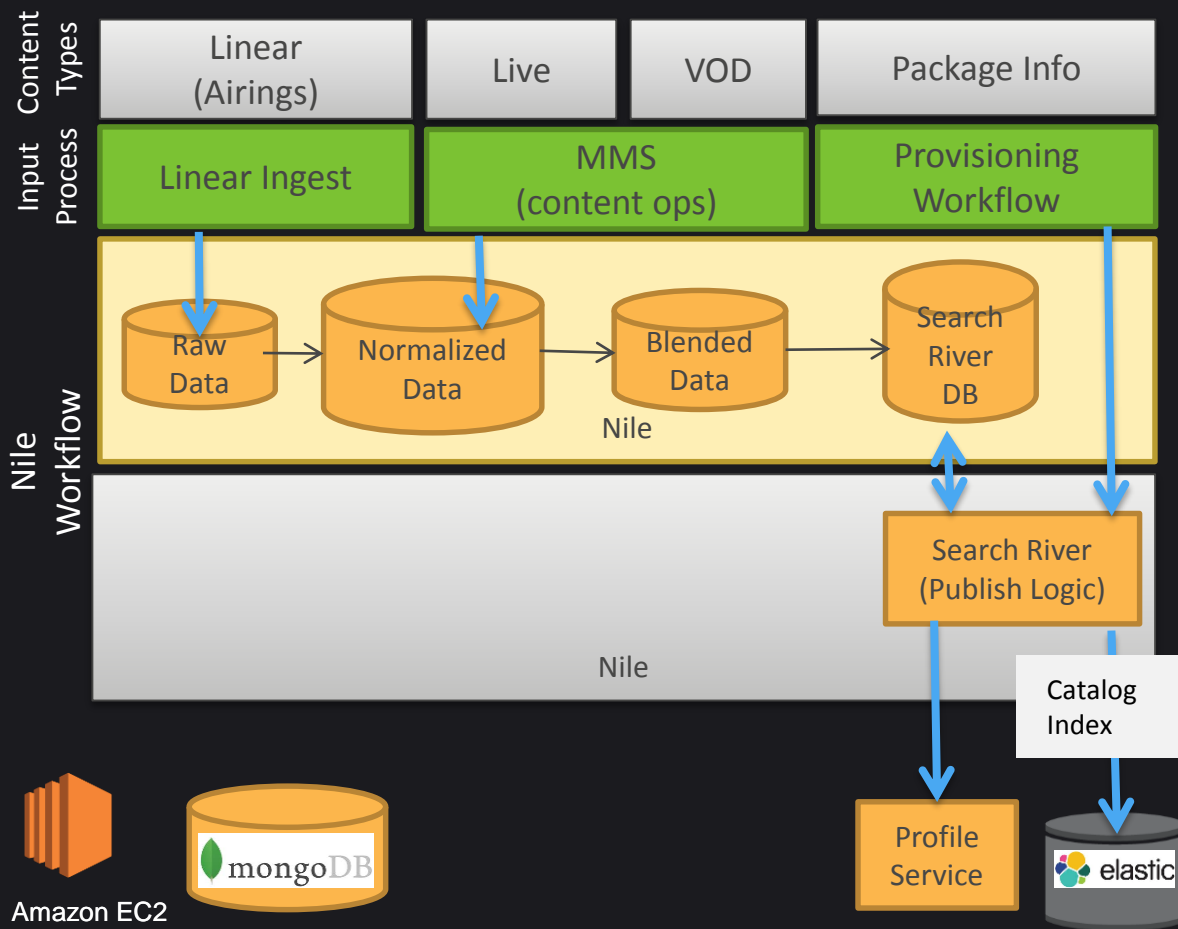


Amazon EC2

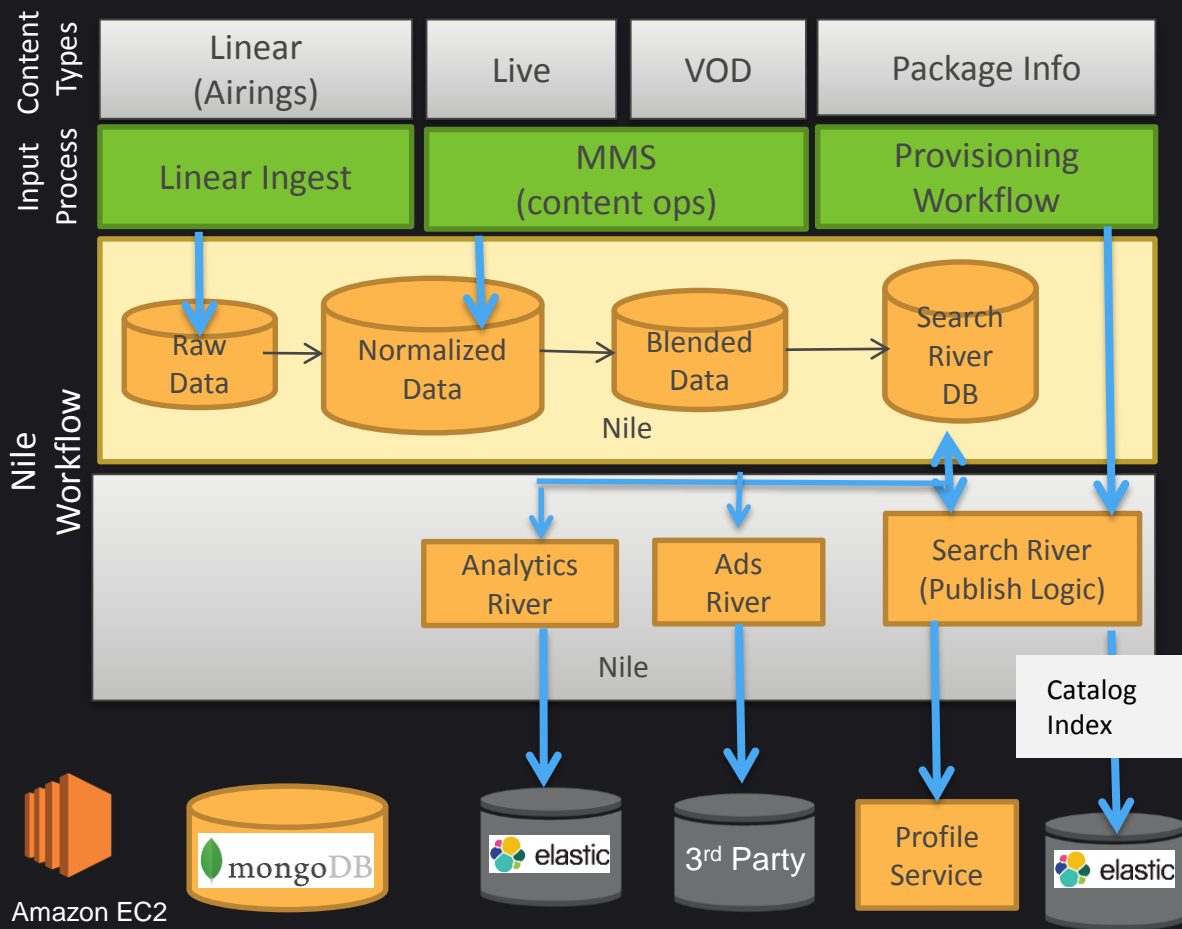


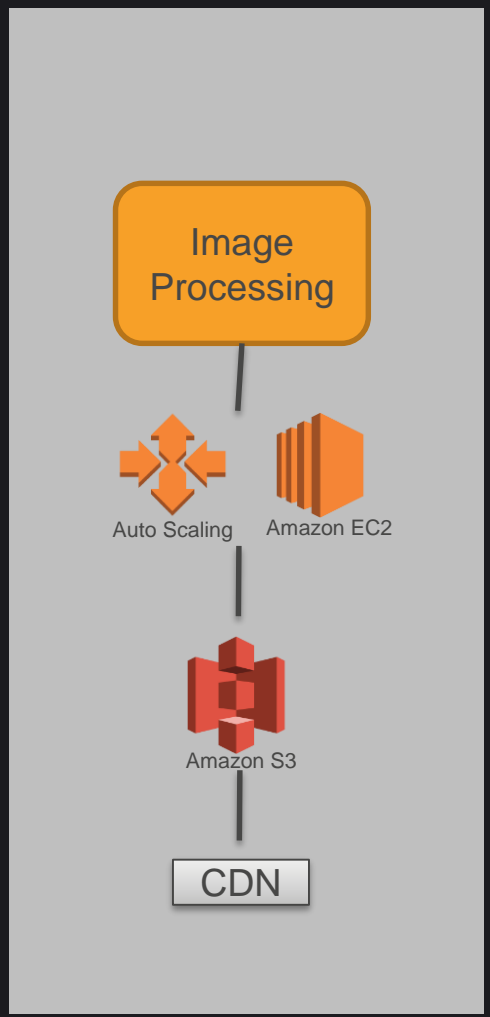
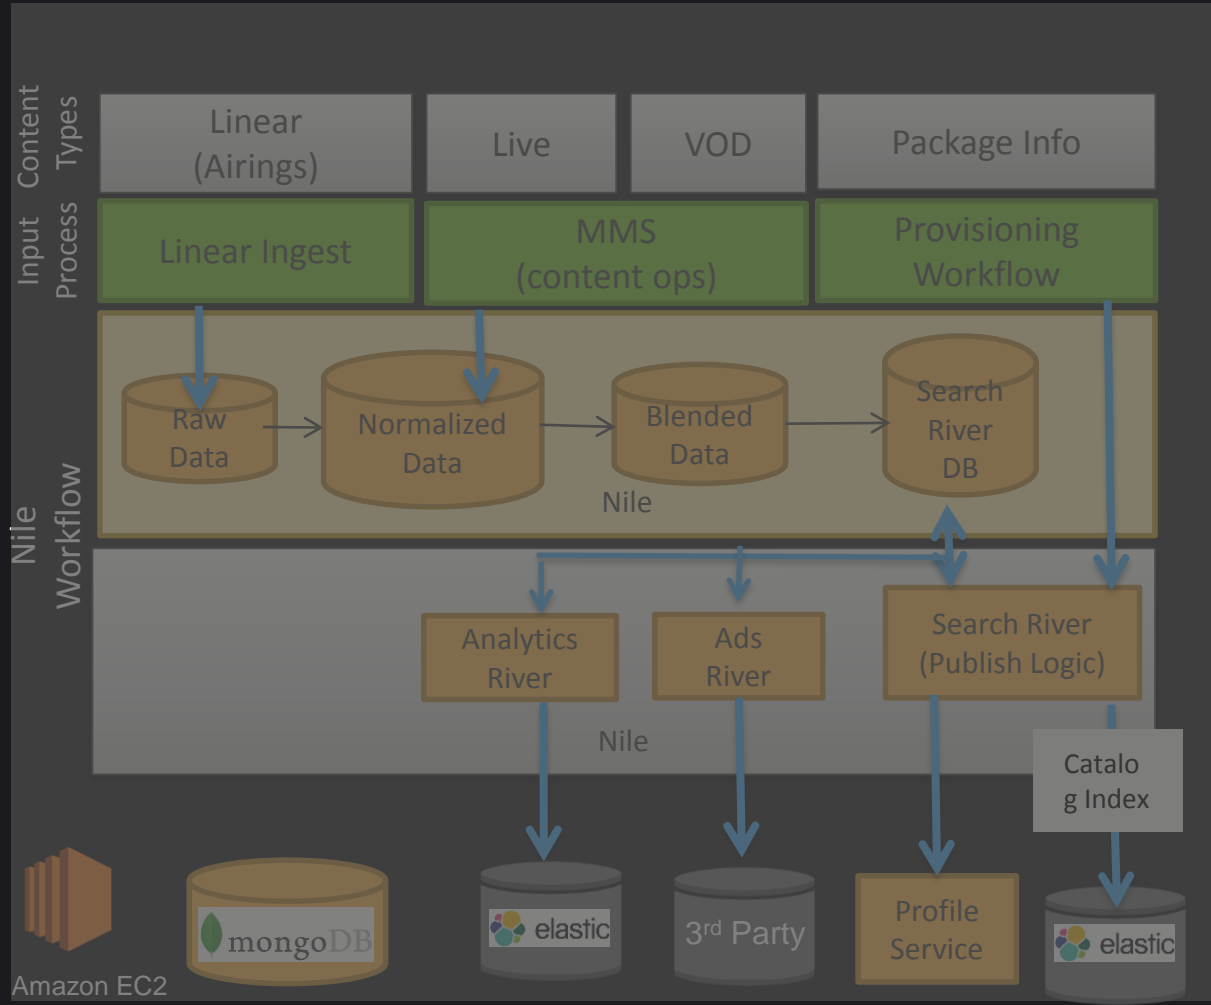
Amazon EC2














# Product Intelligence

## Data Collection



Camus  
(Consume events)



Amazon S3

## Batch Processing



Amazon  
S3



Amazon  
EMR




Amazon  
SWF



## Real-time Processing



Amazon  
S3



## Data Warehouse



Amazon Redshift





# Best Practices

Avoid a priori optimization

Deploy often

Keep simple and separate systems

Don't always give in to “right tool for the right job”



# Summary

# Summary

- Avoid a priori optimization
- Deploy often
- Keep simple and separate systems
- Don't always give in to “right tool for the right job”



# Summary

- Serverless by default
- Avoid lock-in
- Be cautious of the outside
- Reuse data for good
- Avoid a priori optimization
- Deploy often
- Keep simple and separate systems
- Don't always give in to "right tool for the right job"





# Summary

- Serverless by default
- Avoid lock-in
- Be cautious of the outside
- Reuse data for good
- Avoid a priori optimization
- Deploy often
- Keep simple and separate systems
- Don't always give in to "right tool for the right job"

Go Build.

The logo for AWS re:Invent. 'AWS' is in white, 're:' is in orange, and 'Invent' is in white. The background is a dark blue/black field with a complex pattern of thin, glowing blue and orange lines forming geometric shapes and paths.

**AWS**  
**re:Invent**

**Thank you!**



**Remember to complete  
your evaluations!**