

EyeEm

@rinofm

EyeEm

How we scaled our deep learning infrastructure to process 1000 photos / second while cutting our AWS bill by 10%

Rino Montiel, VP Engineering

@rinofm

A person with a vibrant pink wig and a bright red coat stands on a beach, holding a vintage-style camera with a flash. The background shows the ocean and a cloudy sky. The word "EyeEm" is overlaid in large white text across the center of the image.

EyeEm

A person with short, wavy pink hair and a bright red coat is holding a vintage-style camera with a large external flash. They are standing on a beach with the ocean in the background. The image has a semi-transparent dark overlay.

EyeEm

EyeEm is a photography company that builds the world's leading computer vision technology to connect its global creative community with iconic brands.

EyeEm

Photo quality rating

Popular

Auto tagging

Model release detection

What we've automated

Keywording

Captioning

Search

Personalized aesthetics

Concept training platform

EyeEm

Understanding photography



EyeEm



Man in Abandoned Building

EyeEm

INDOORS

BRICK WALL

CHEERFUL

LISTENING

SMARTPHONE

YOUNG WOMAN

WIRELESS TECHNOLOGY

HOLDING



EyeEm

INDOORS

BRICK WALL

CHEERFUL

LISTENING

SMARTPHONE


89%

Aesthetic Score

YOUNG WOMAN

WIRELESS TECHNOLOGY

HOLDING



EyeEm

Three Parts

Research, Production, Data



EyeEm

Beginning: Research

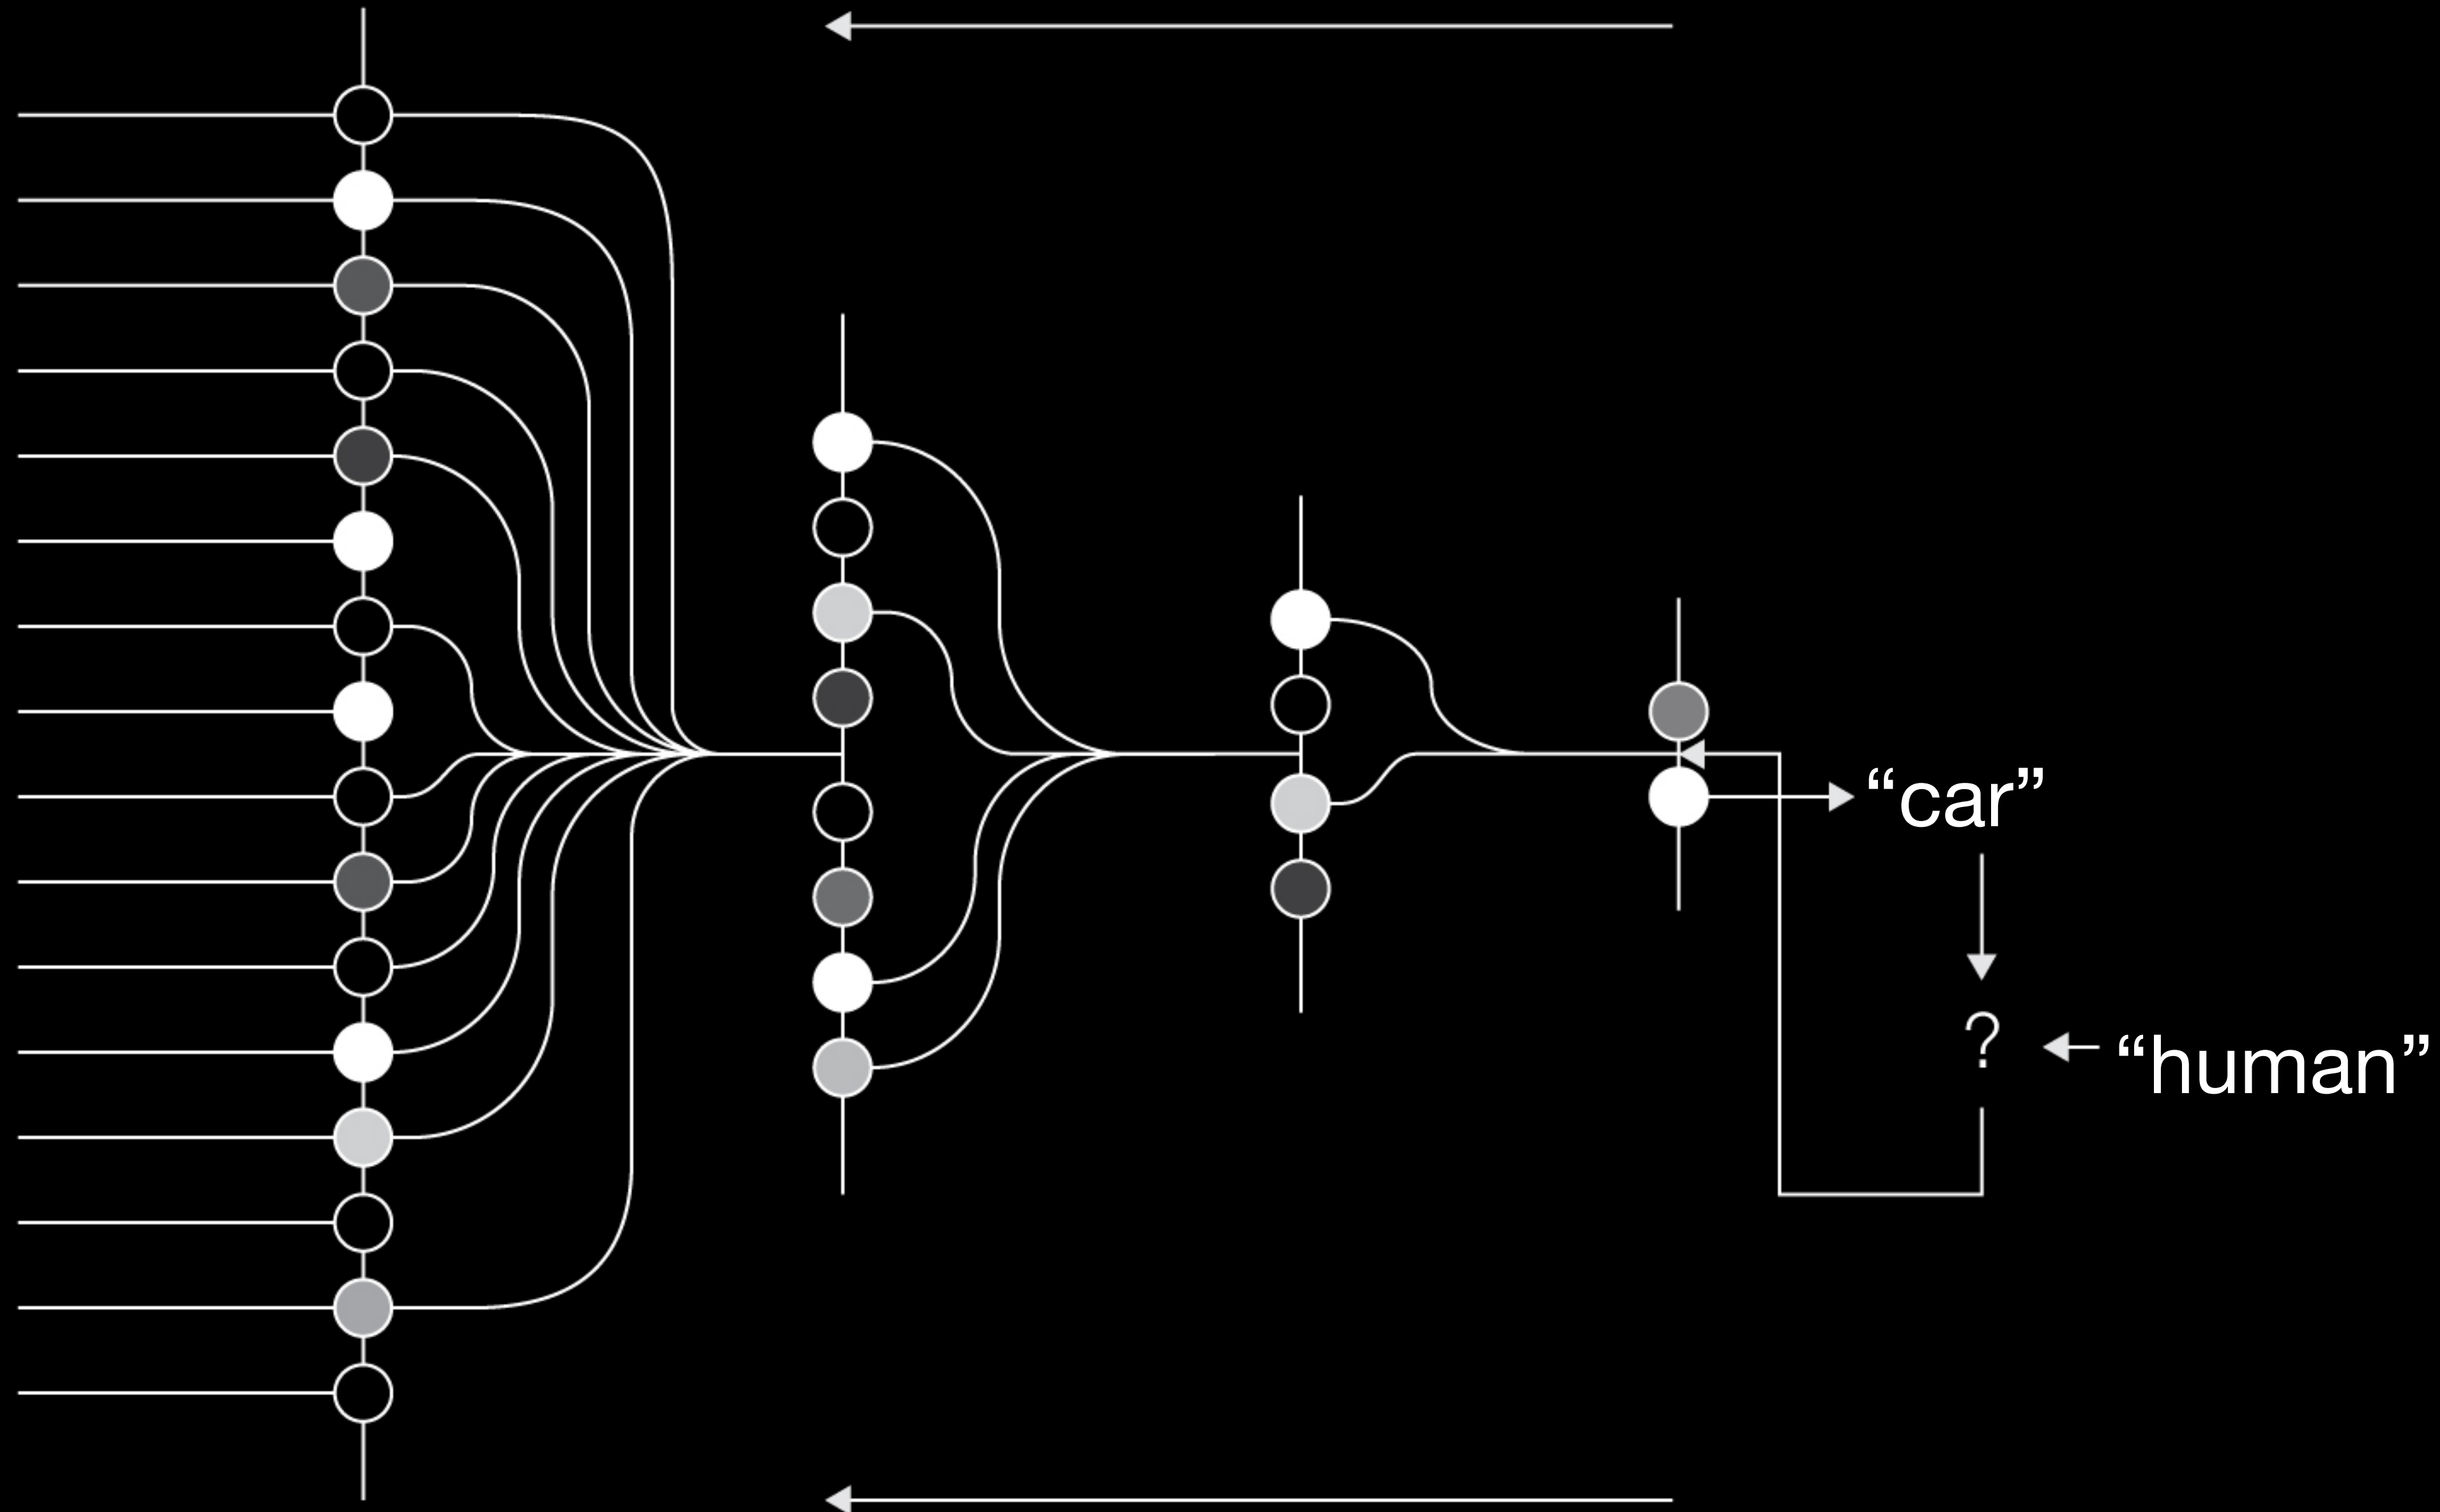

Where the models (and some trouble) come from




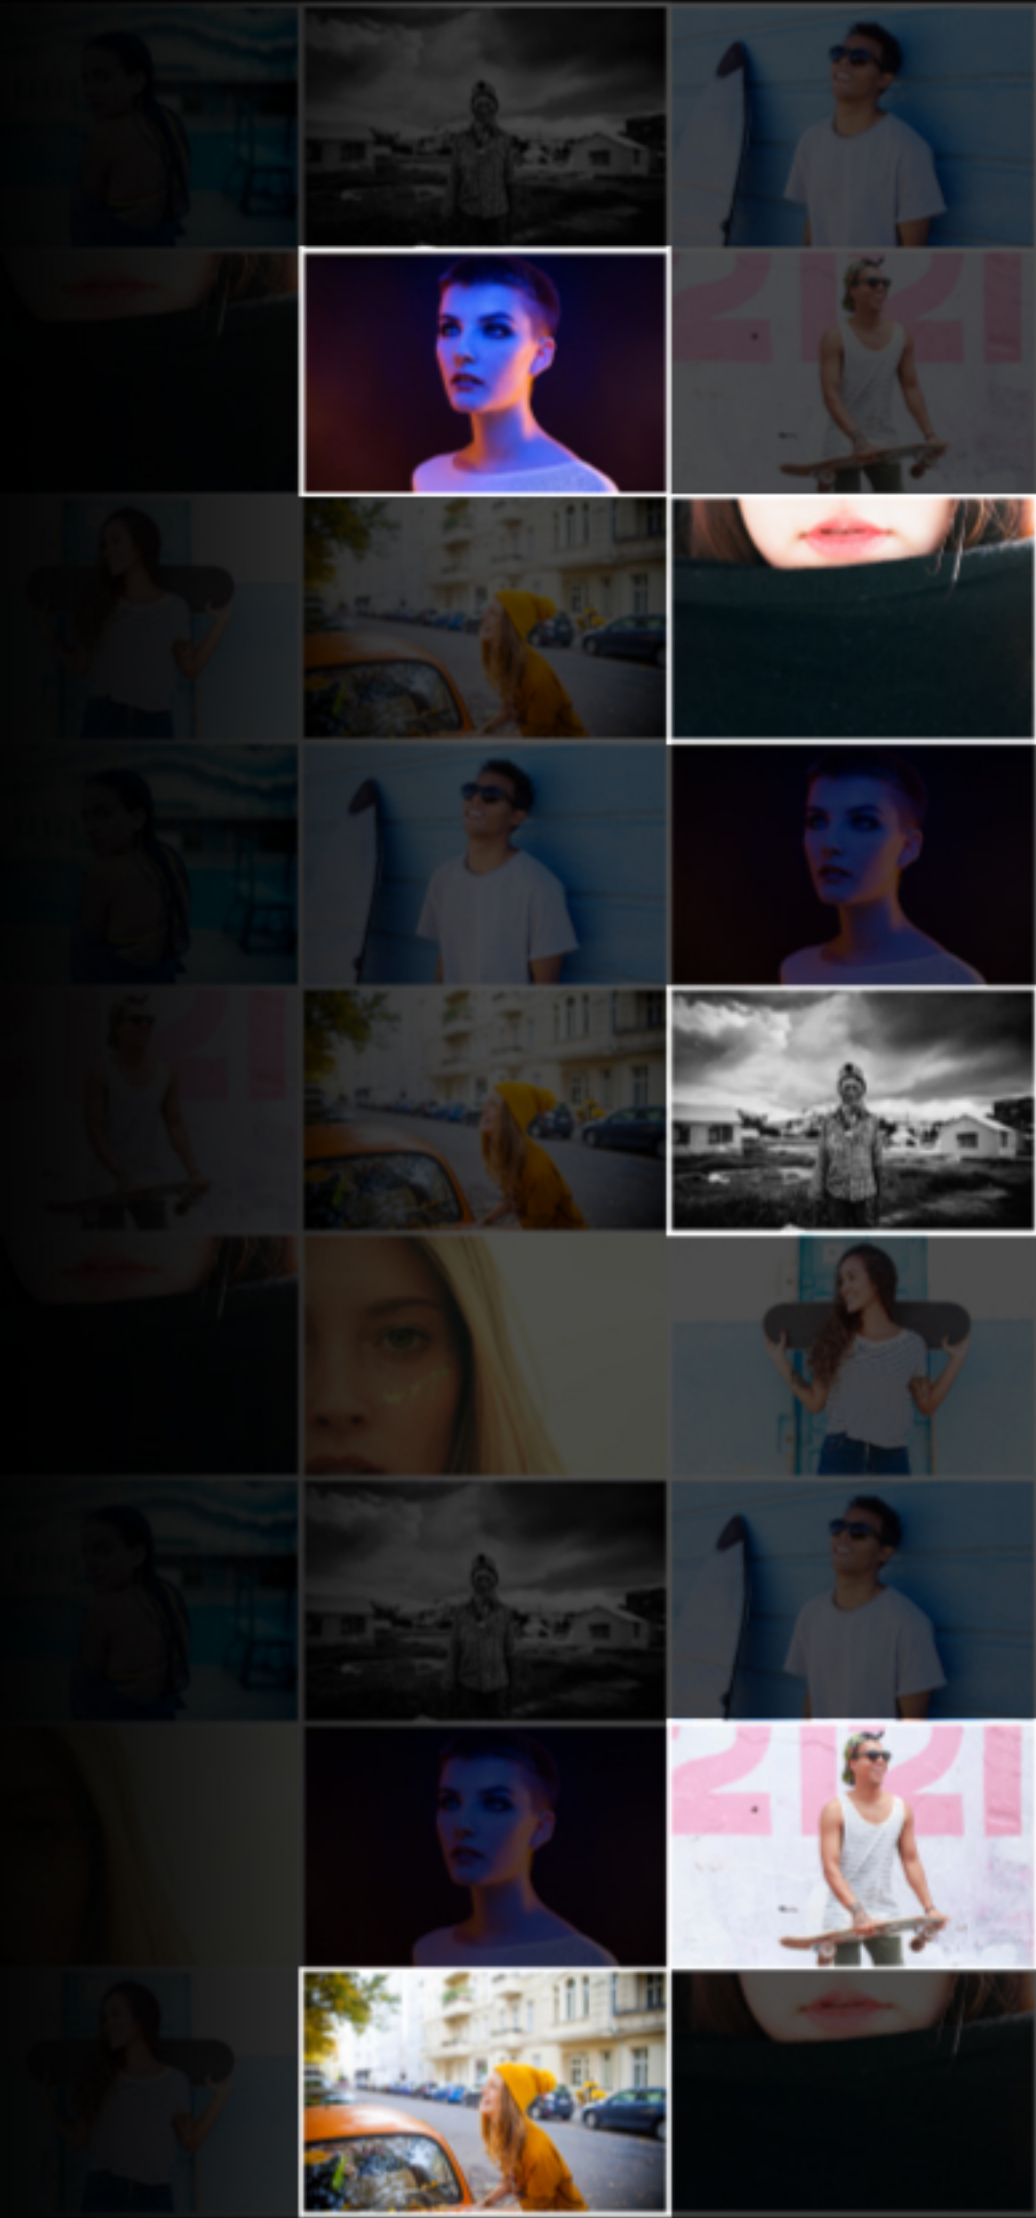
EyeEm

Training and inference

EyeEm



EyeEm



EyeEm

Deep learning and GPUs


GPUs run deep learning training
and inference very efficiently

EyeEm

Training with millions of
photos needs GPUs

More data, more GPUs

EyeEm



EyeEm

Model trained and
production ready

EyeEm

S3 - Original photo



Thumbnail Service



Deep Learning Service

Keywords
Aesthetics
Captions

EyeEm

Problem:
Input size changes

EyeEm

Problem:
Resizing algorithm

EyeEm

Problem: Thumbnail Service


Scaling thumbnail generation for training

EyeEm

Problem: Thumbnail Service

Scaling thumbnail generation for training

Thumbnail Service



EyeEm

R&D

Thumbnail Internal

Thumbnail Service



Problem: Thumbnail Service

Scaling thumbnail generation for training

EyeEm

Training: Learning new functionality from past data

Inference: Applying functionality to new data

Input size can change

Resizing algorithm affects performance

Thumbnail generation is load intensive

Rely on a mix of on-demand and spot instances \$

EyeEm

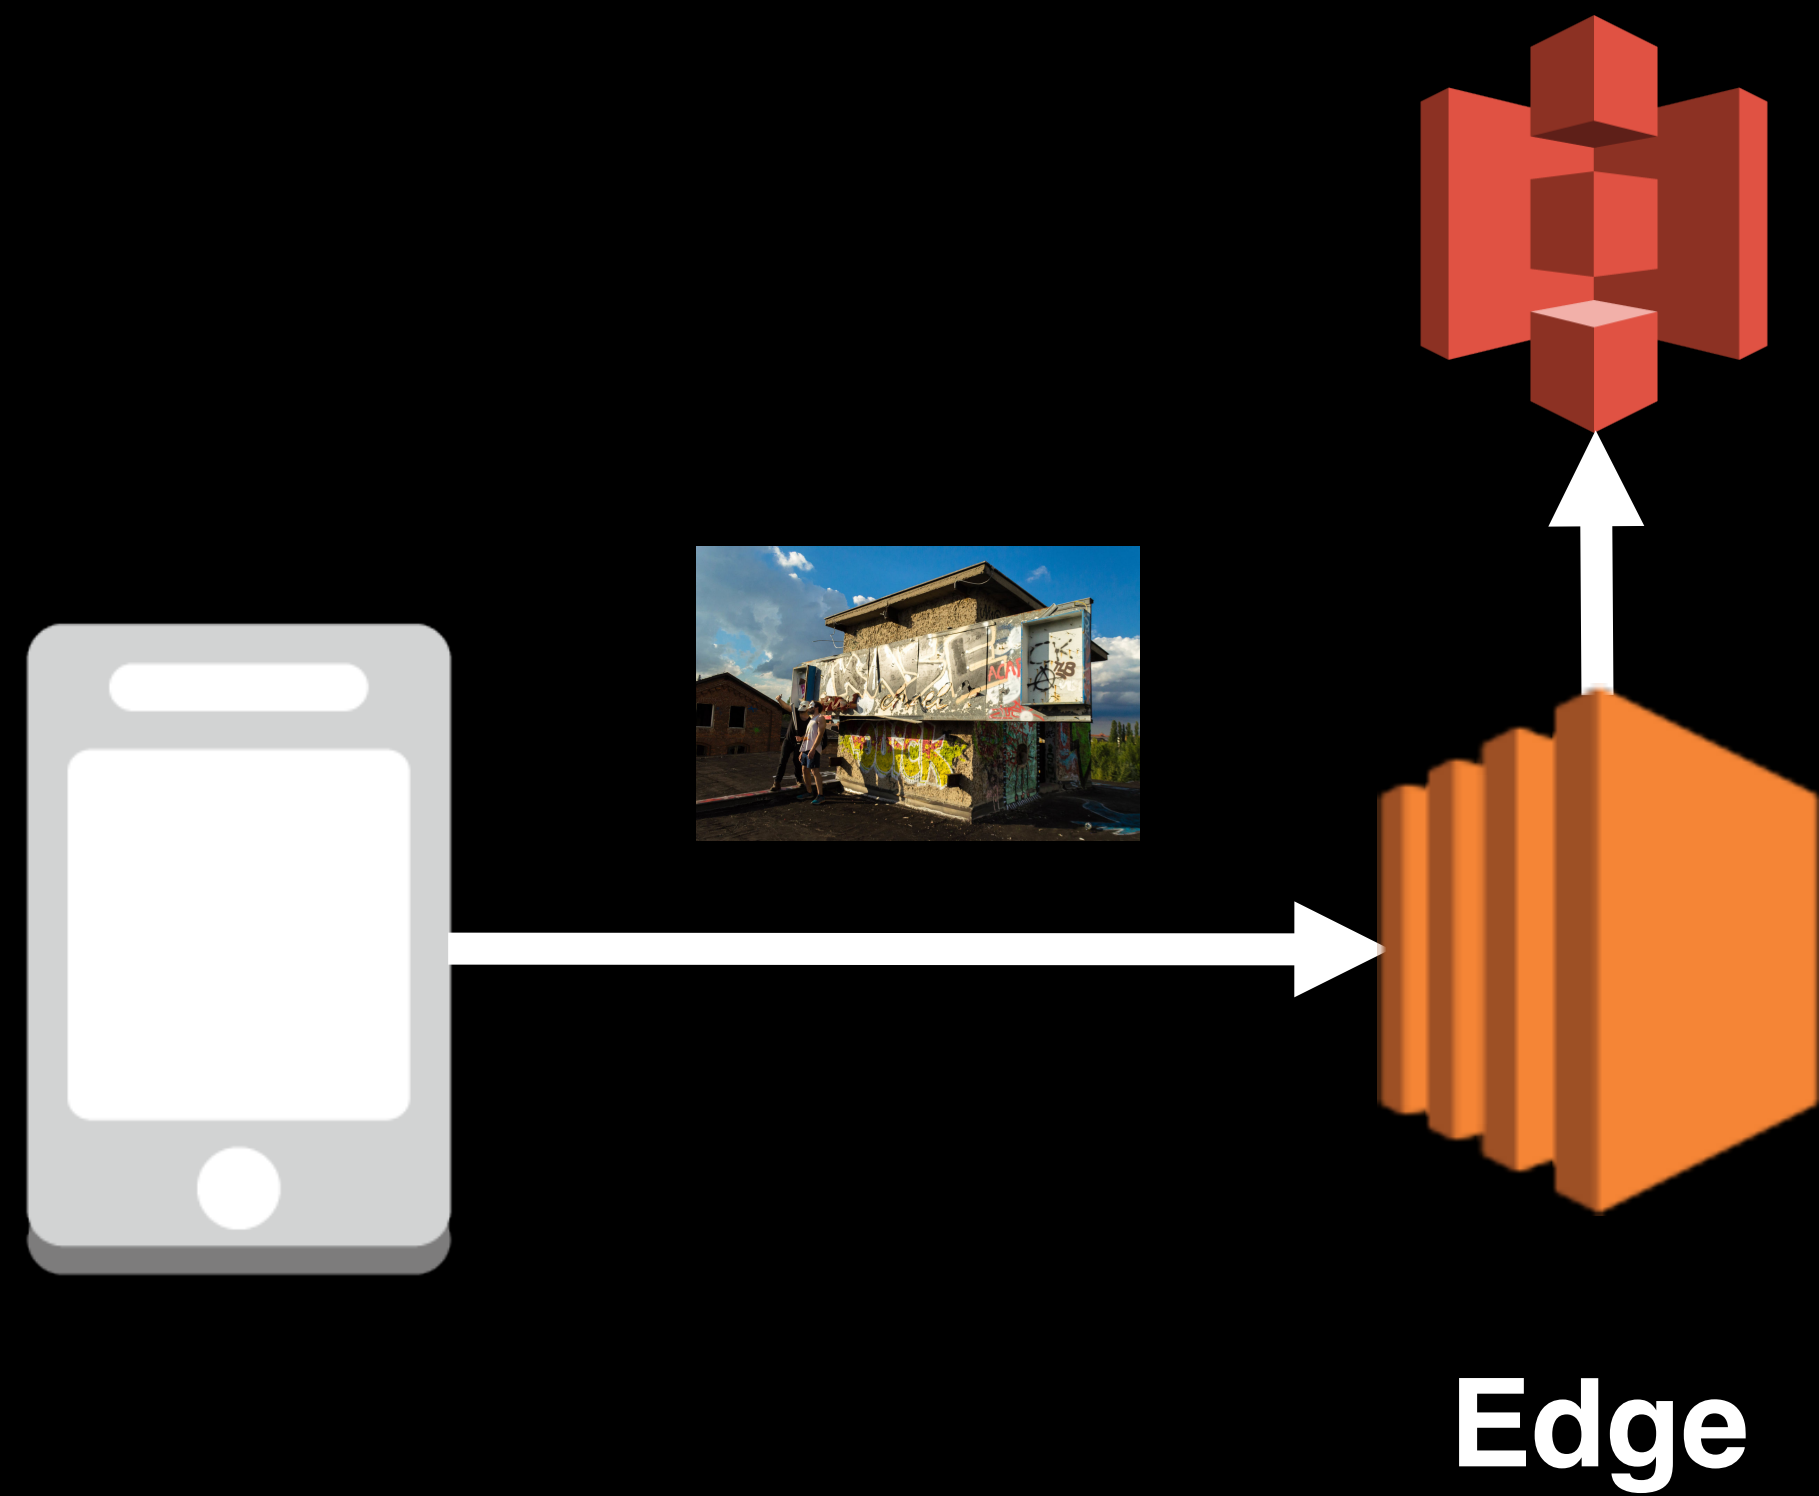
Middle: Production

Dealing with scale



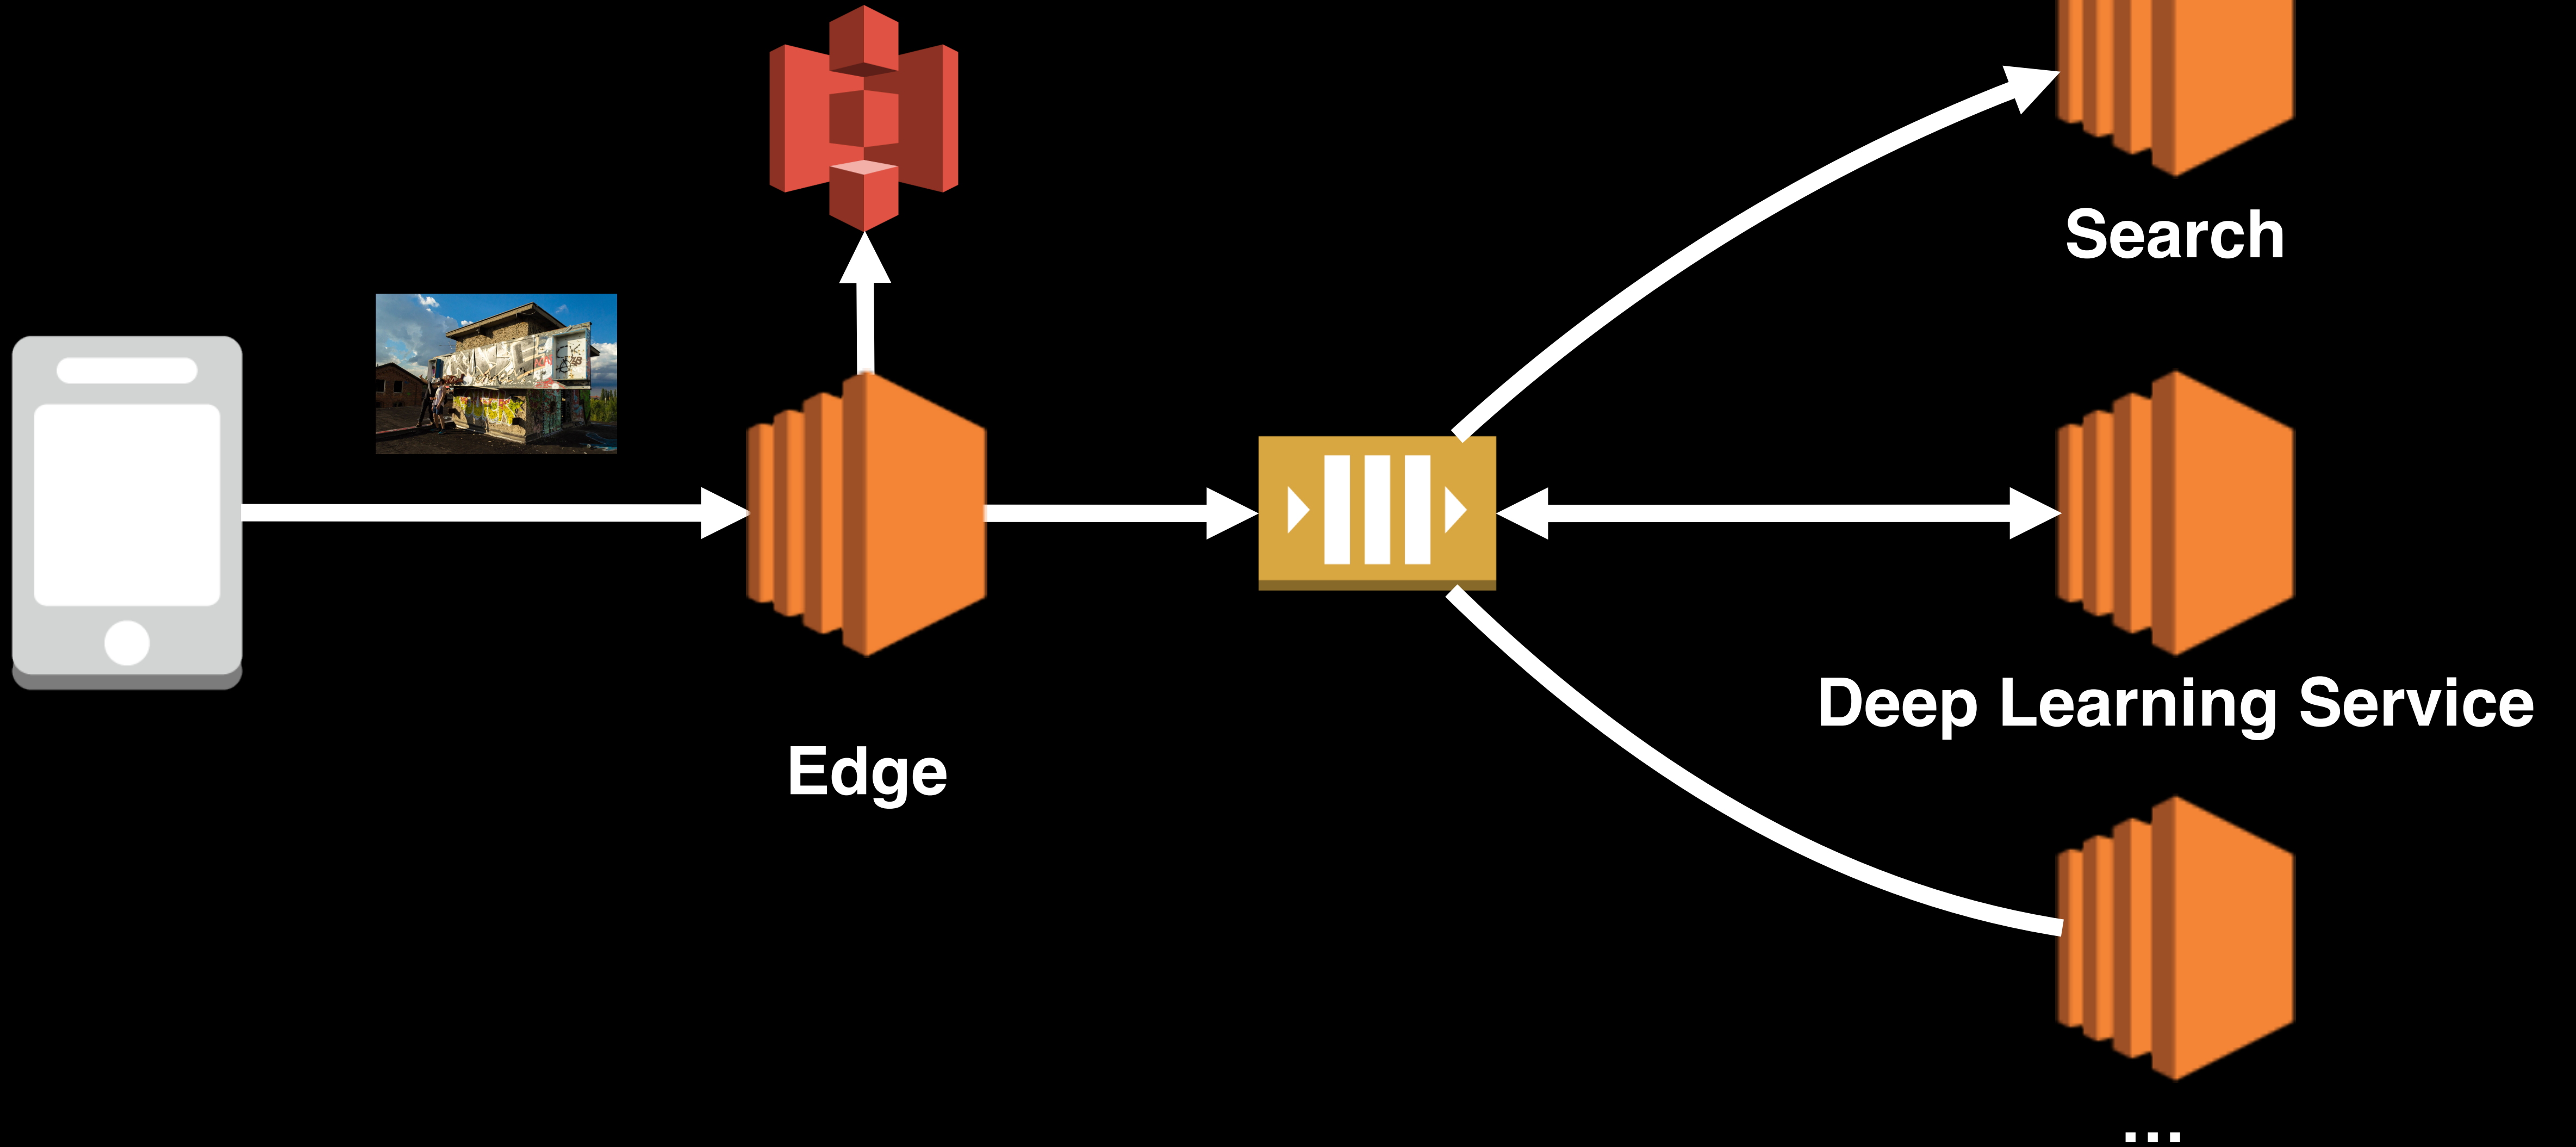
EyeEm

S3 - Original photos



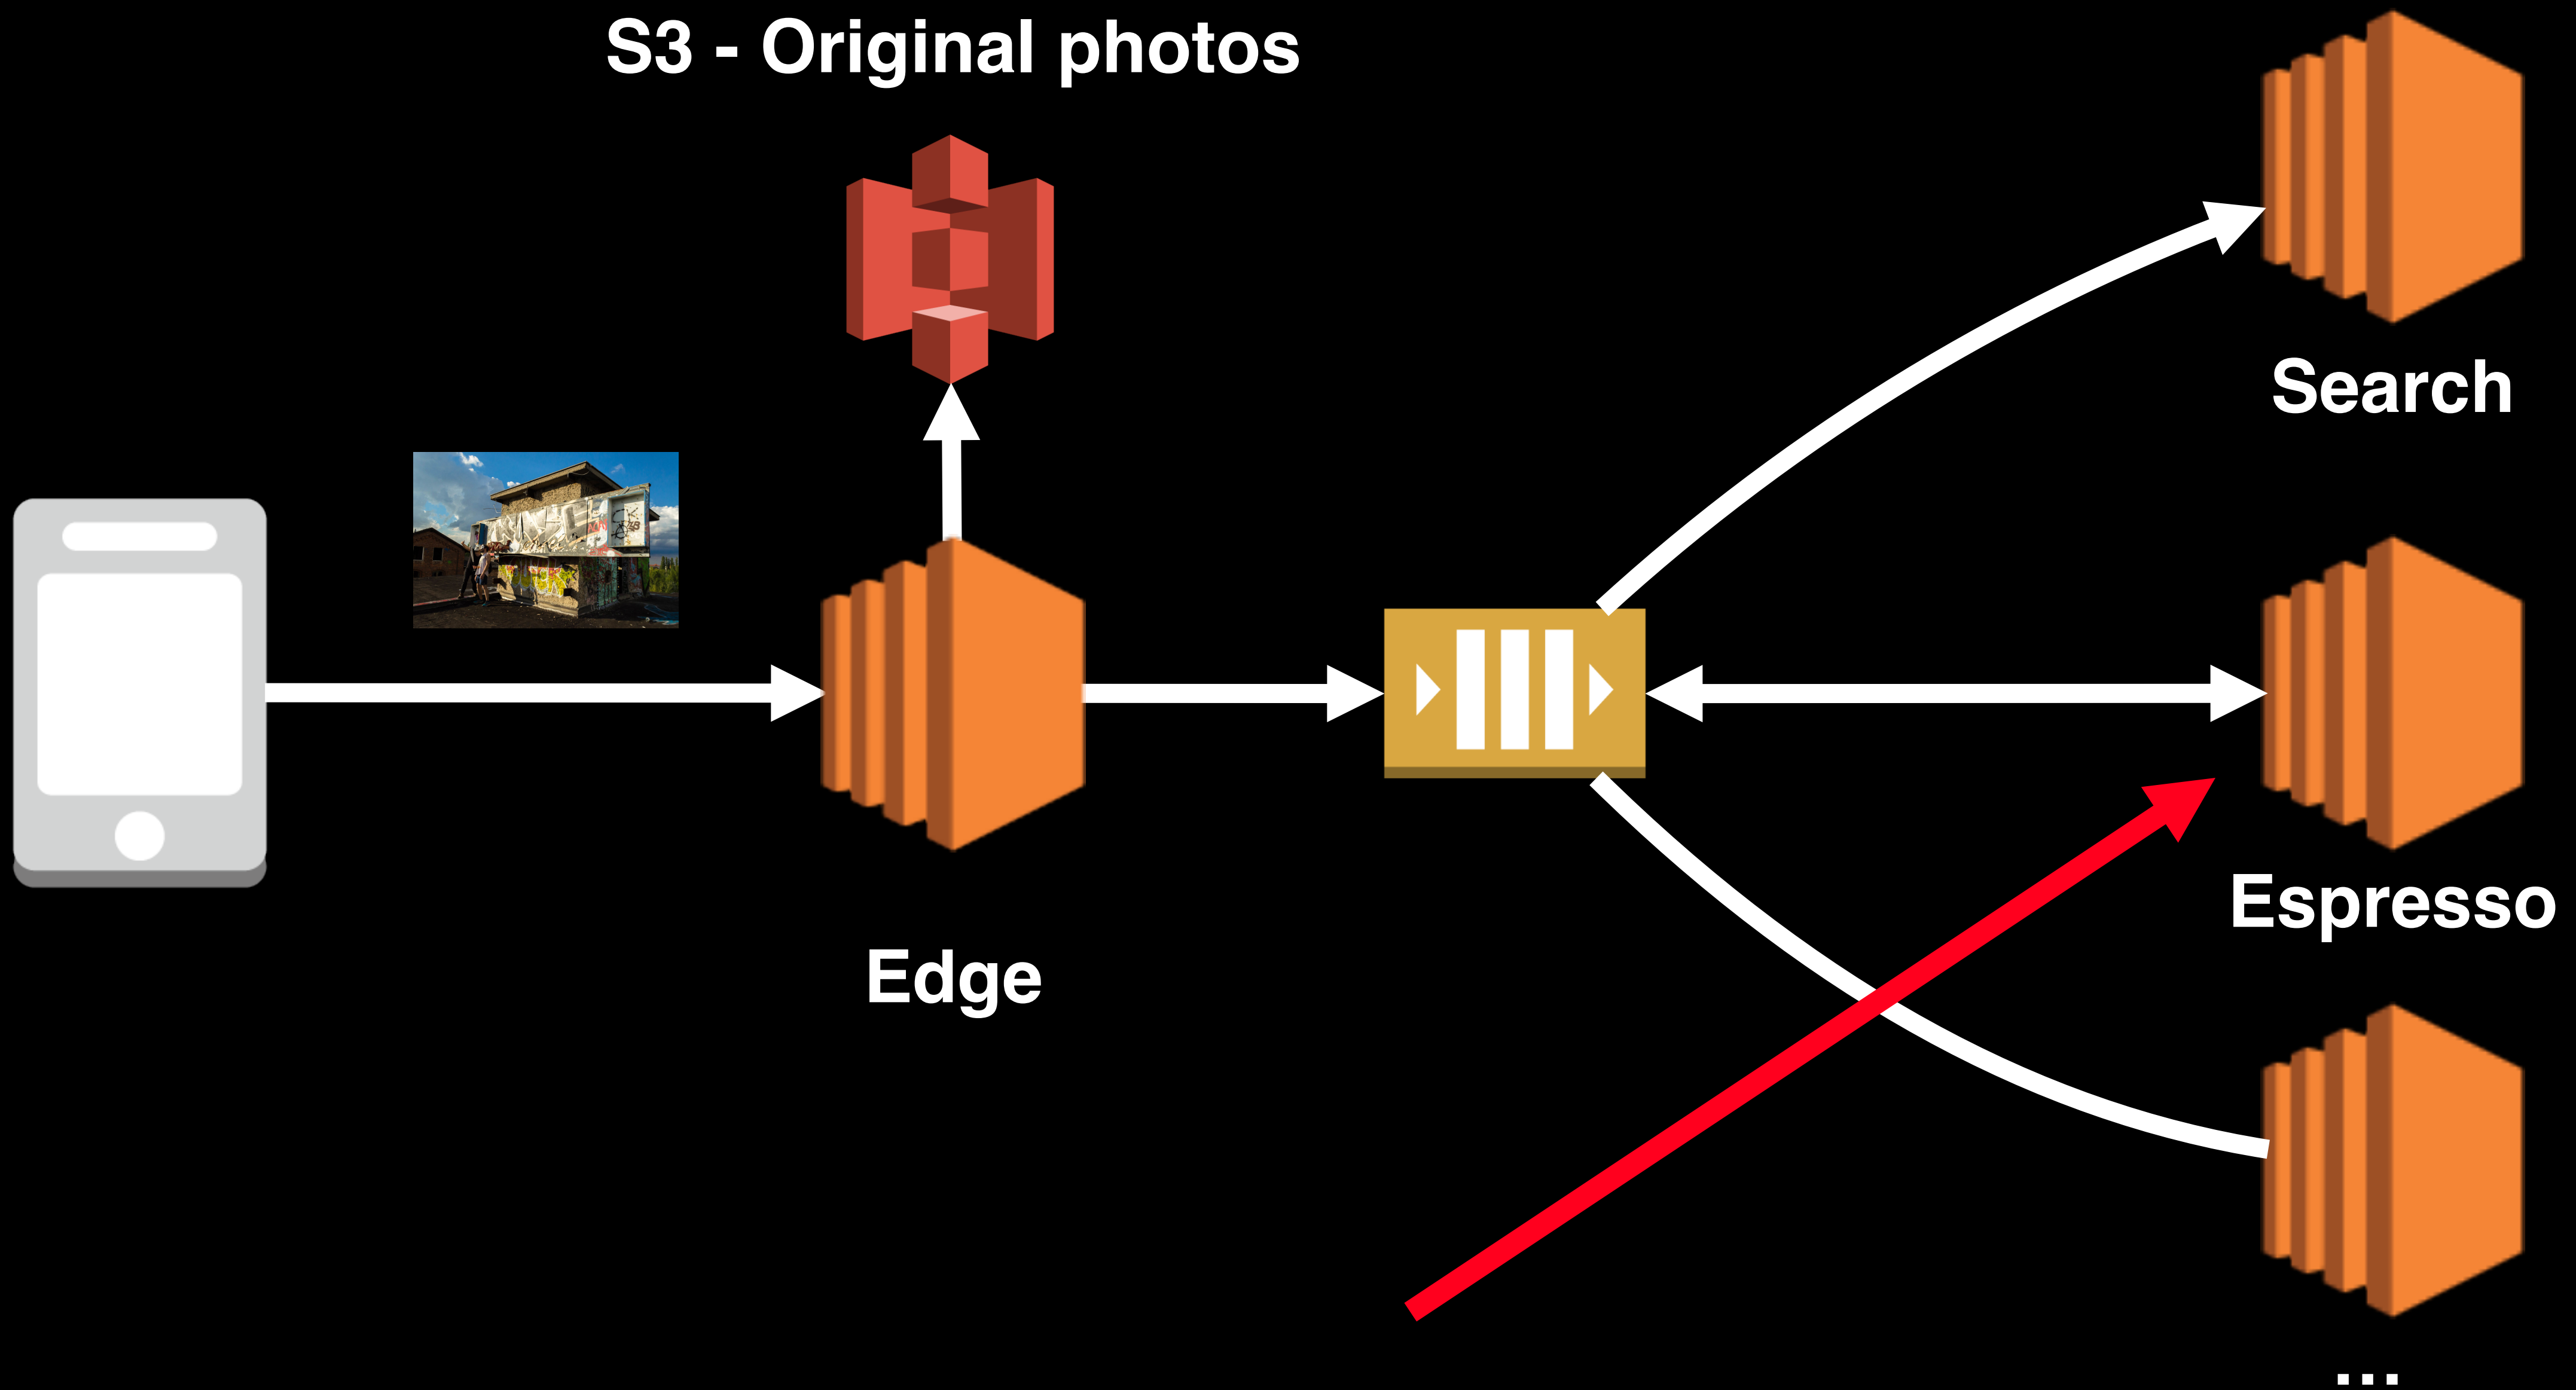
EyeEm

S3 - Original photos



EyeEm

S3 - Original photos



EyeEm

Espresso

G2 machines



aesthetics

concepts

quality

captions

**personalised
aesthetics**

...

EyeEm

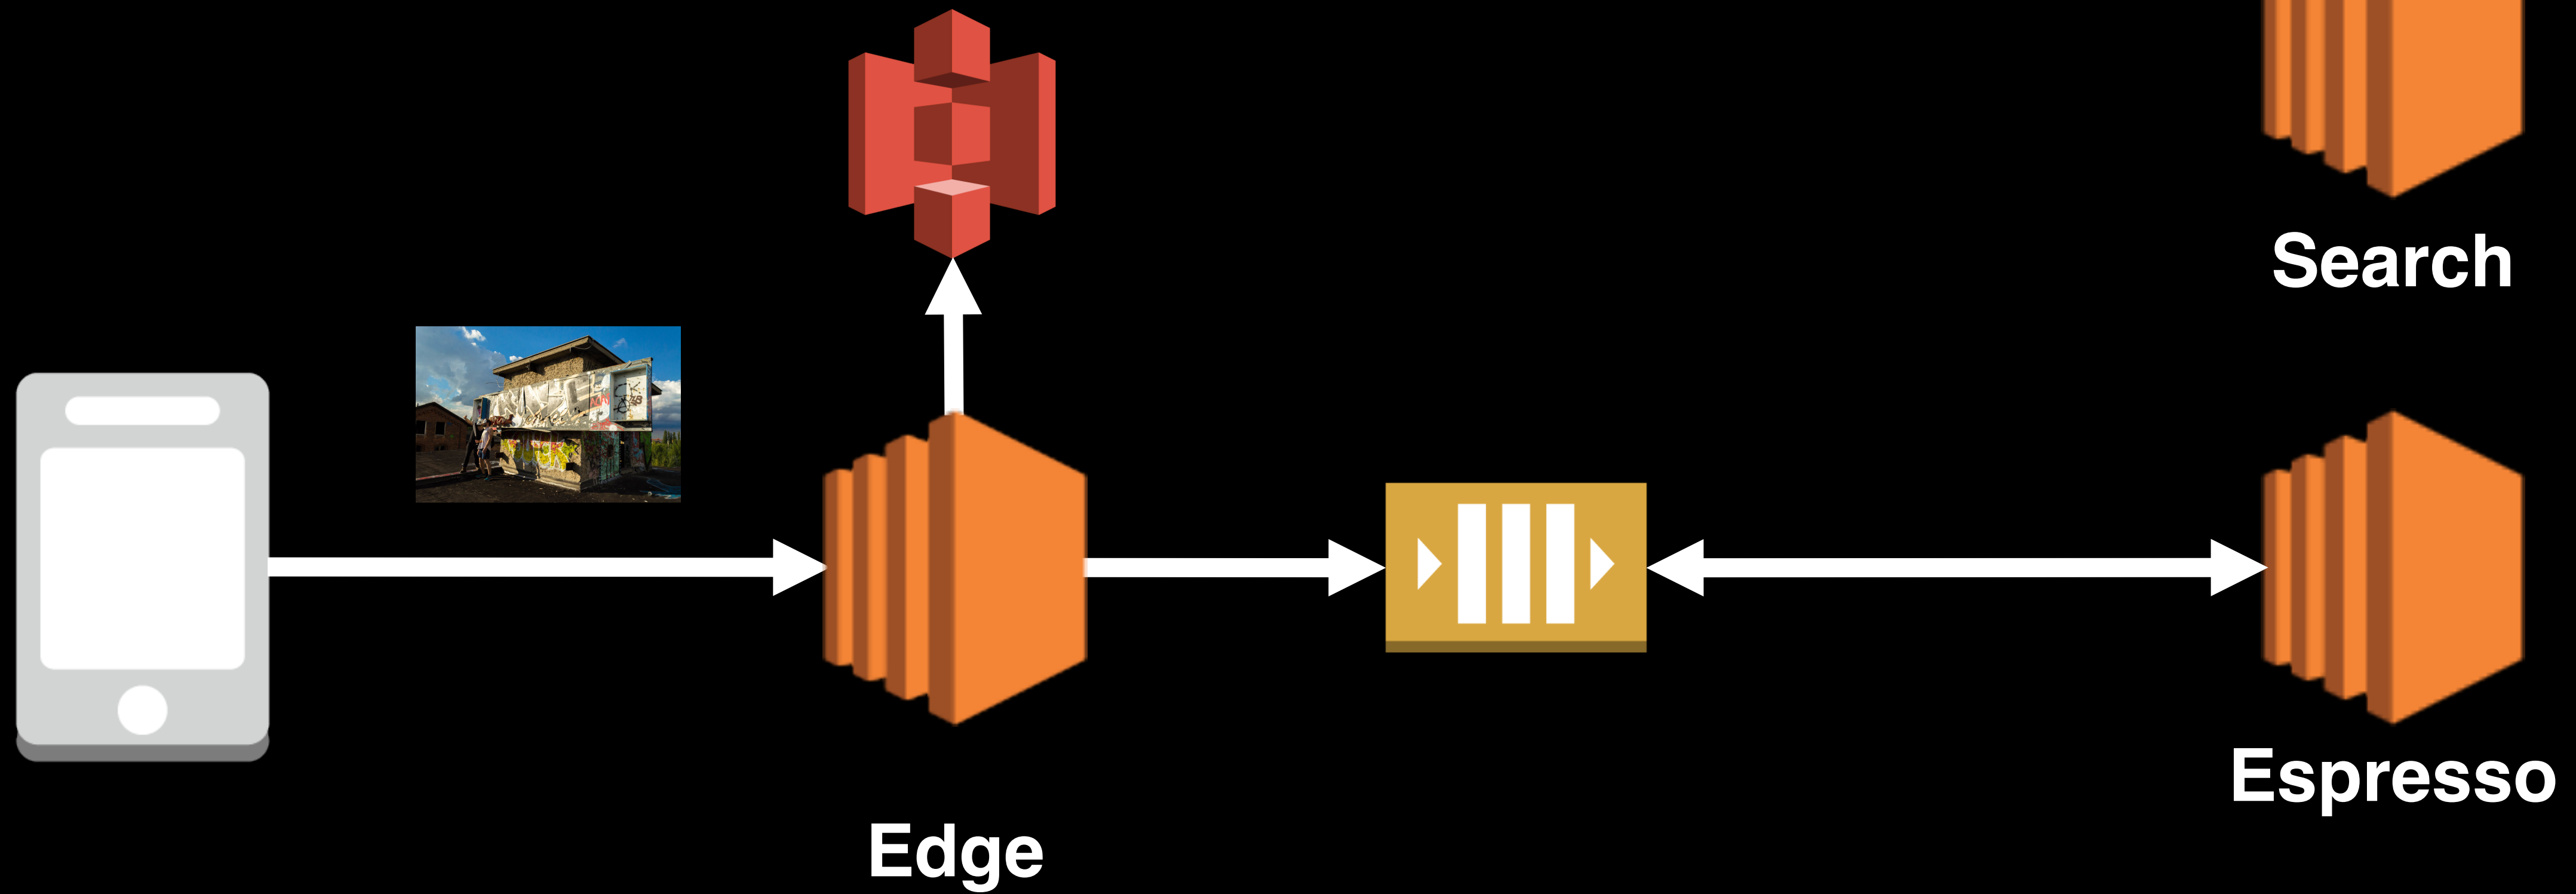
What we've automated

Search



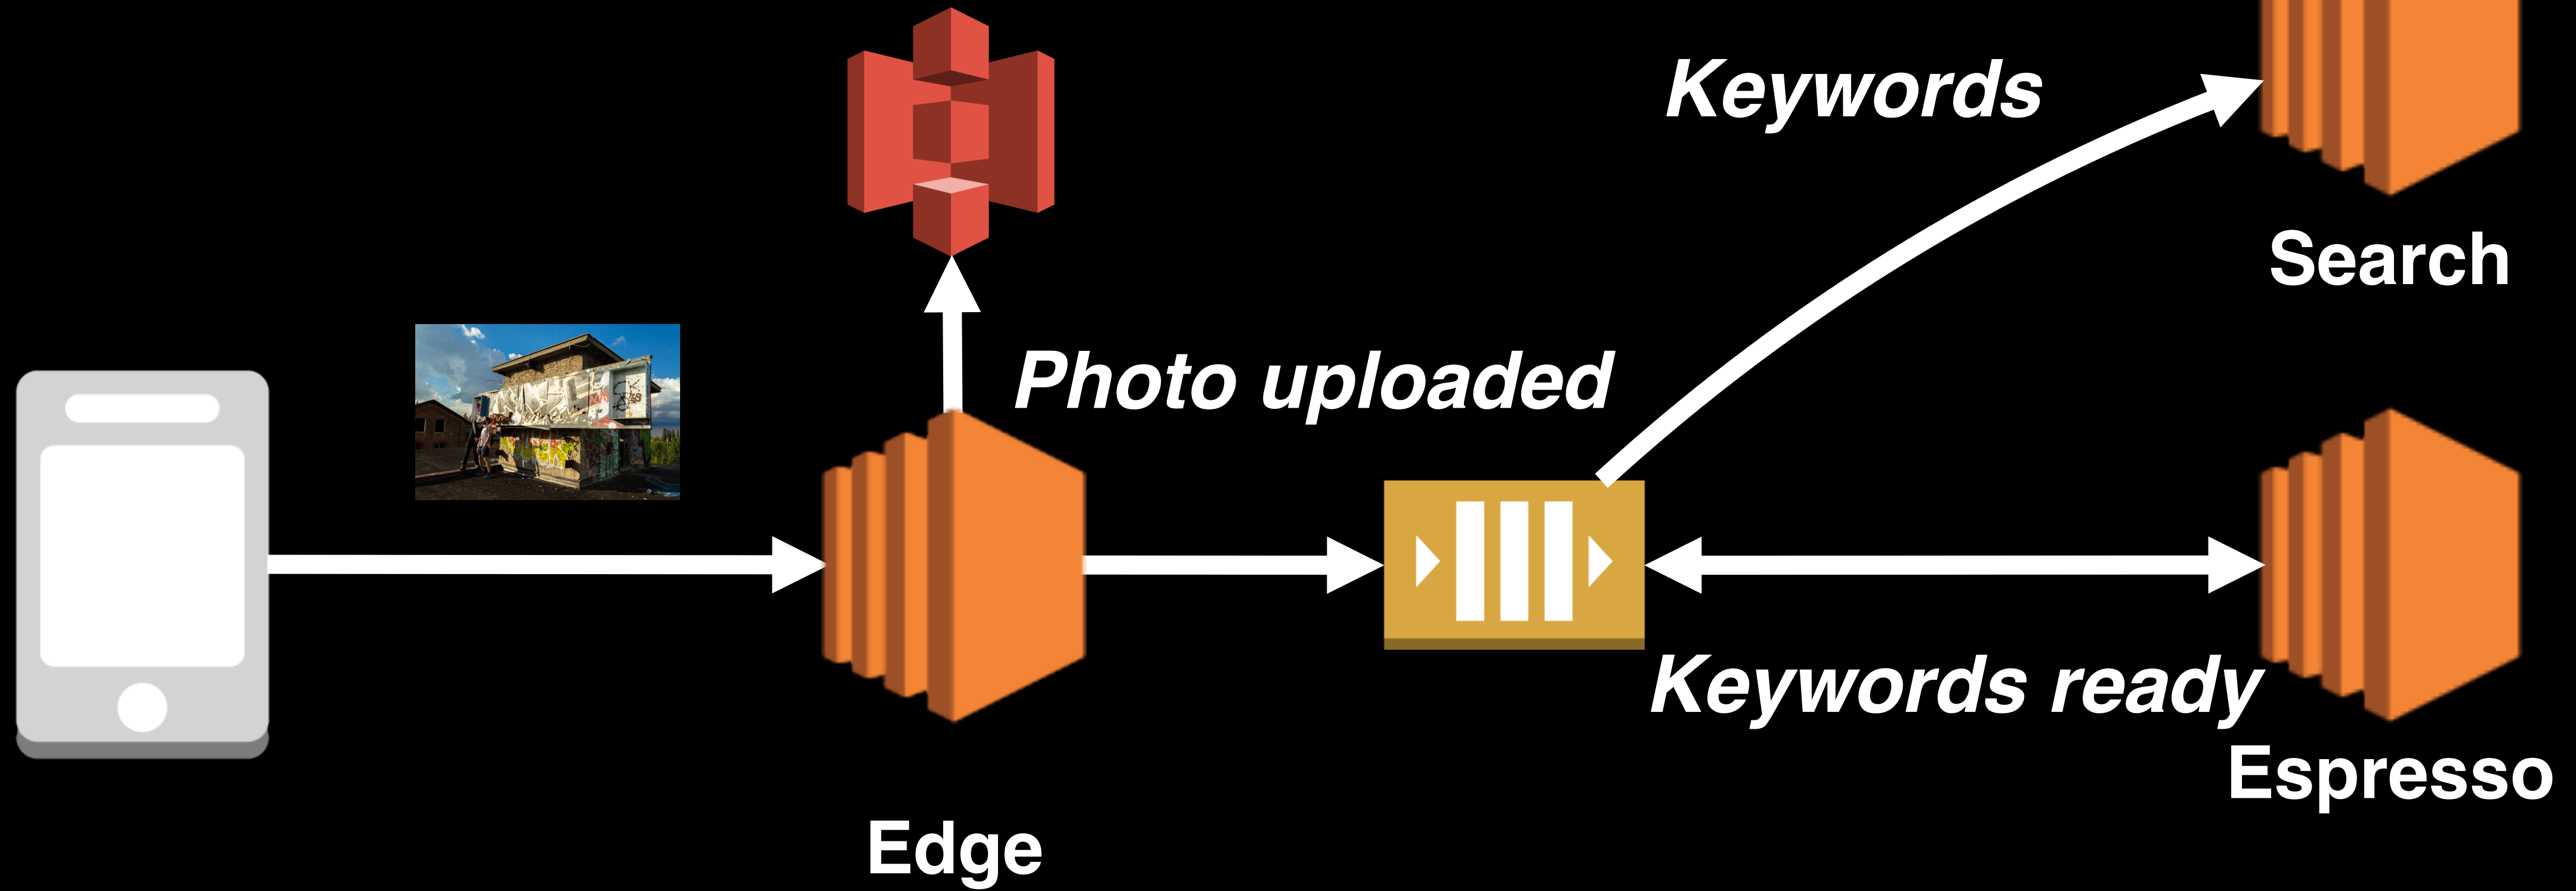
EyeEm

S3 - Original photos



EyeEm

S3 - Original photos



Hide filters —

153,612 results

Sort by Most Relevant

Market —

Photos on EyeEm Market

Format —

- Portrait
- Landscape
- Square
- Widescreen
- Panorama





Size —

- Large
- Medium
- Small





Need help?

Our photo curators can help you find what you're looking for - free of charge.







 Erik Nardini
about one year ...  3866  112 







 Madeline Yates
11 months ago  35  




 Giri Mohan
6 months ago  24  

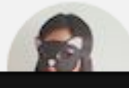



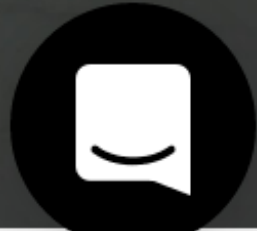


 Uri Peleg
over one year ago  10  



 Nasty Disaronko 

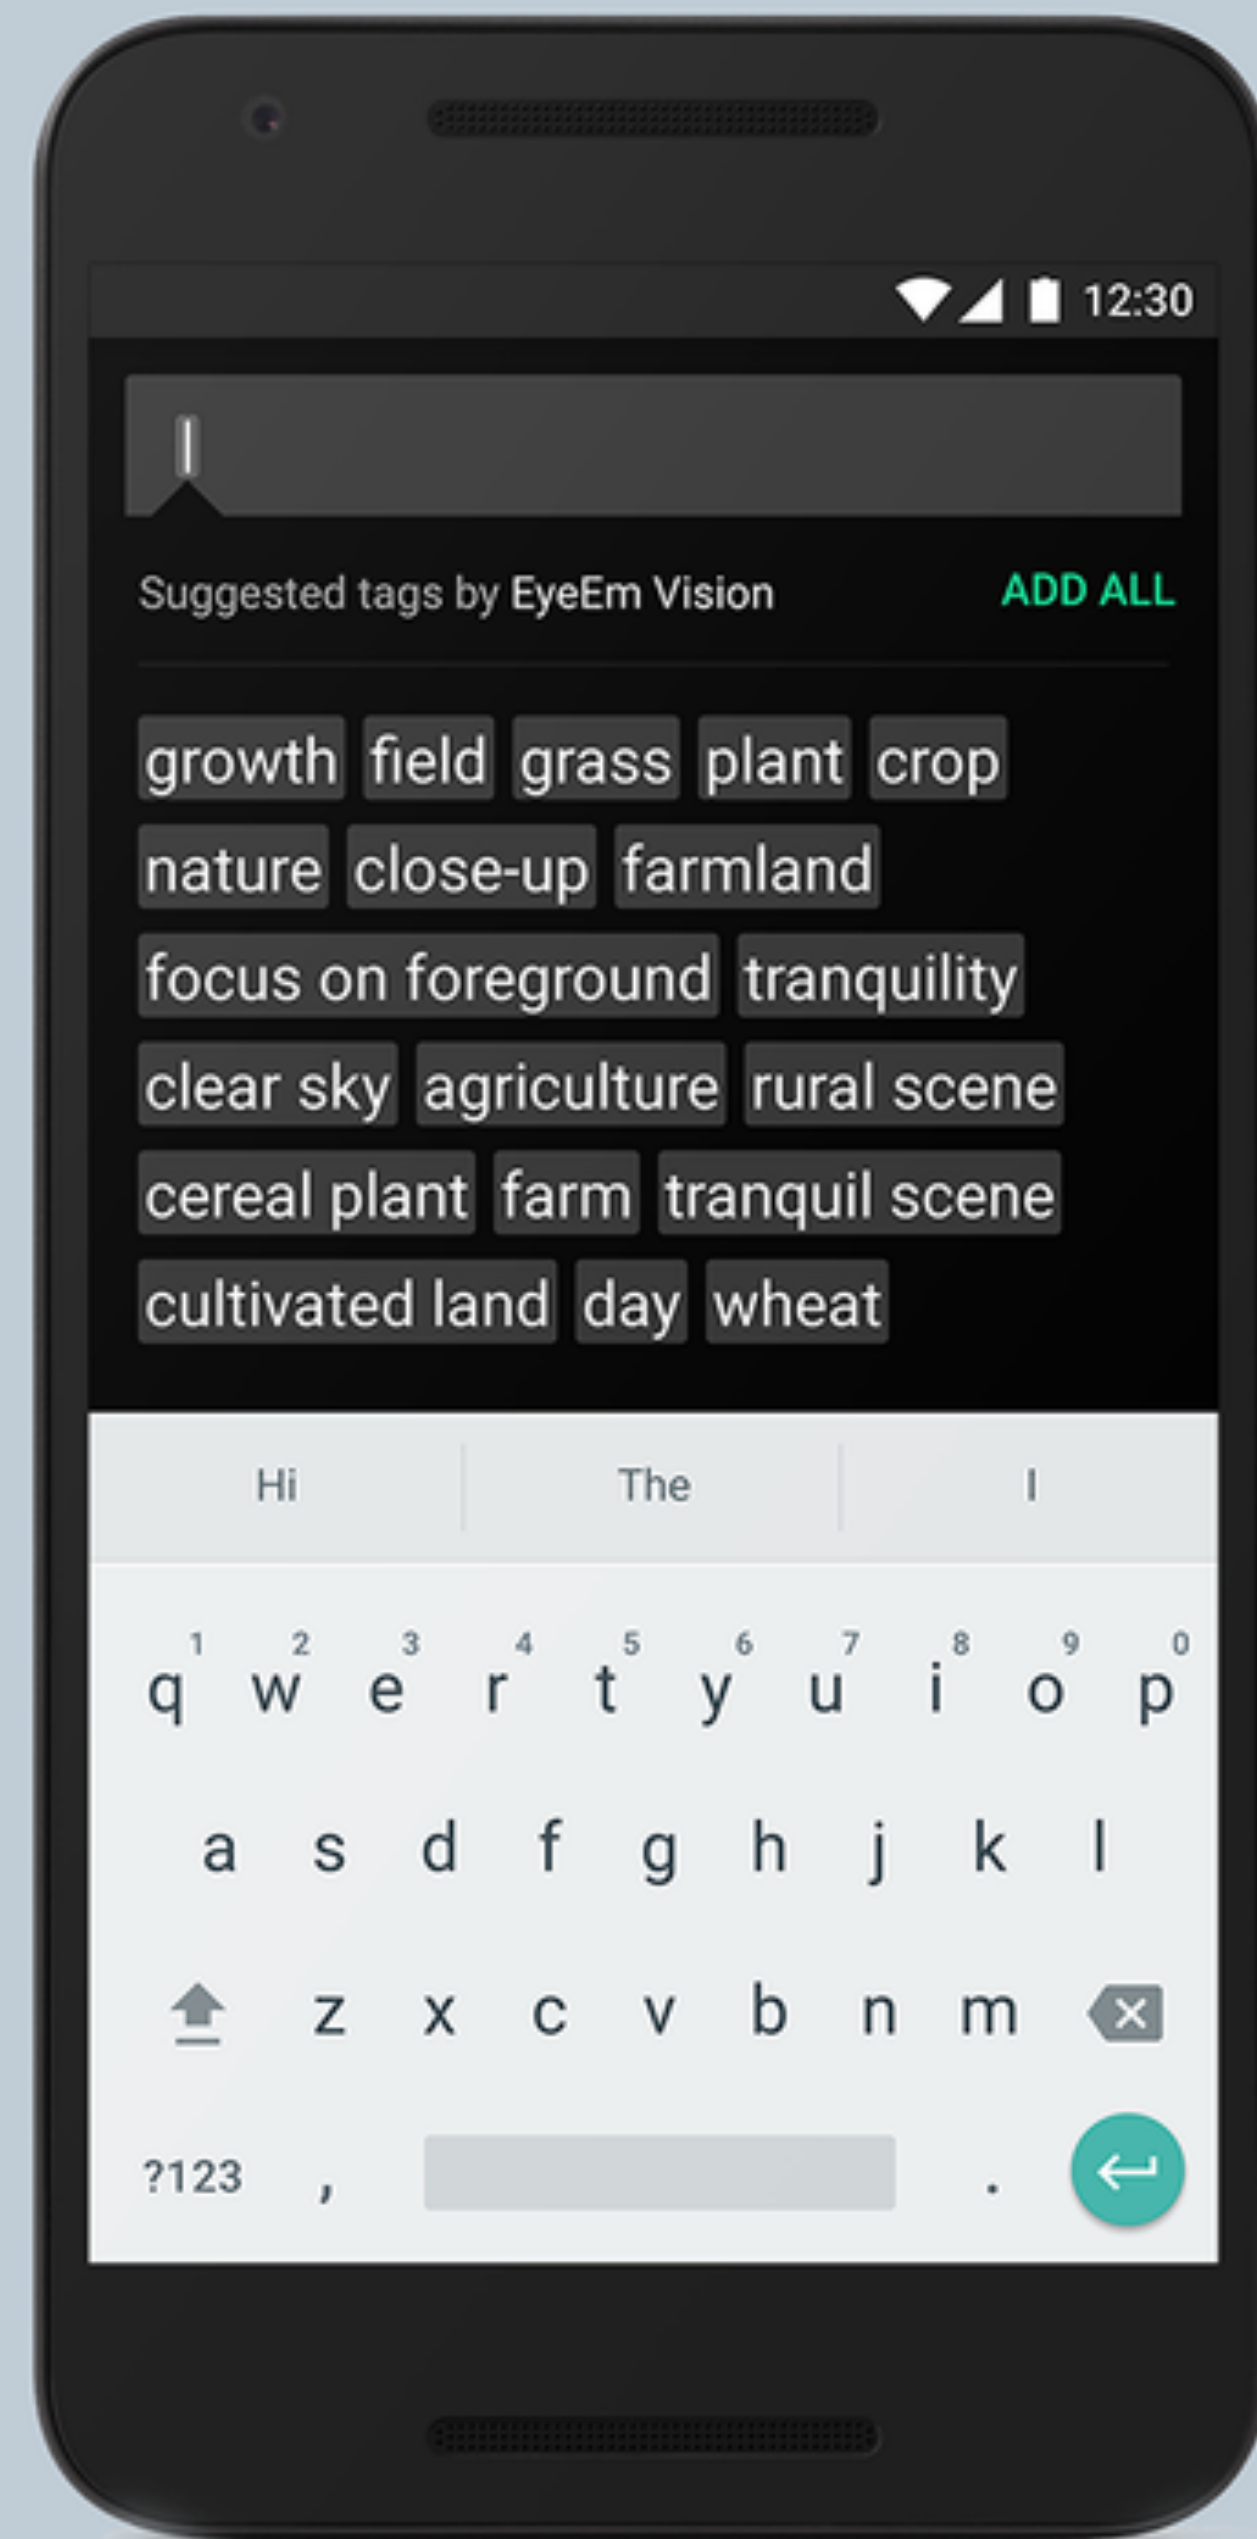
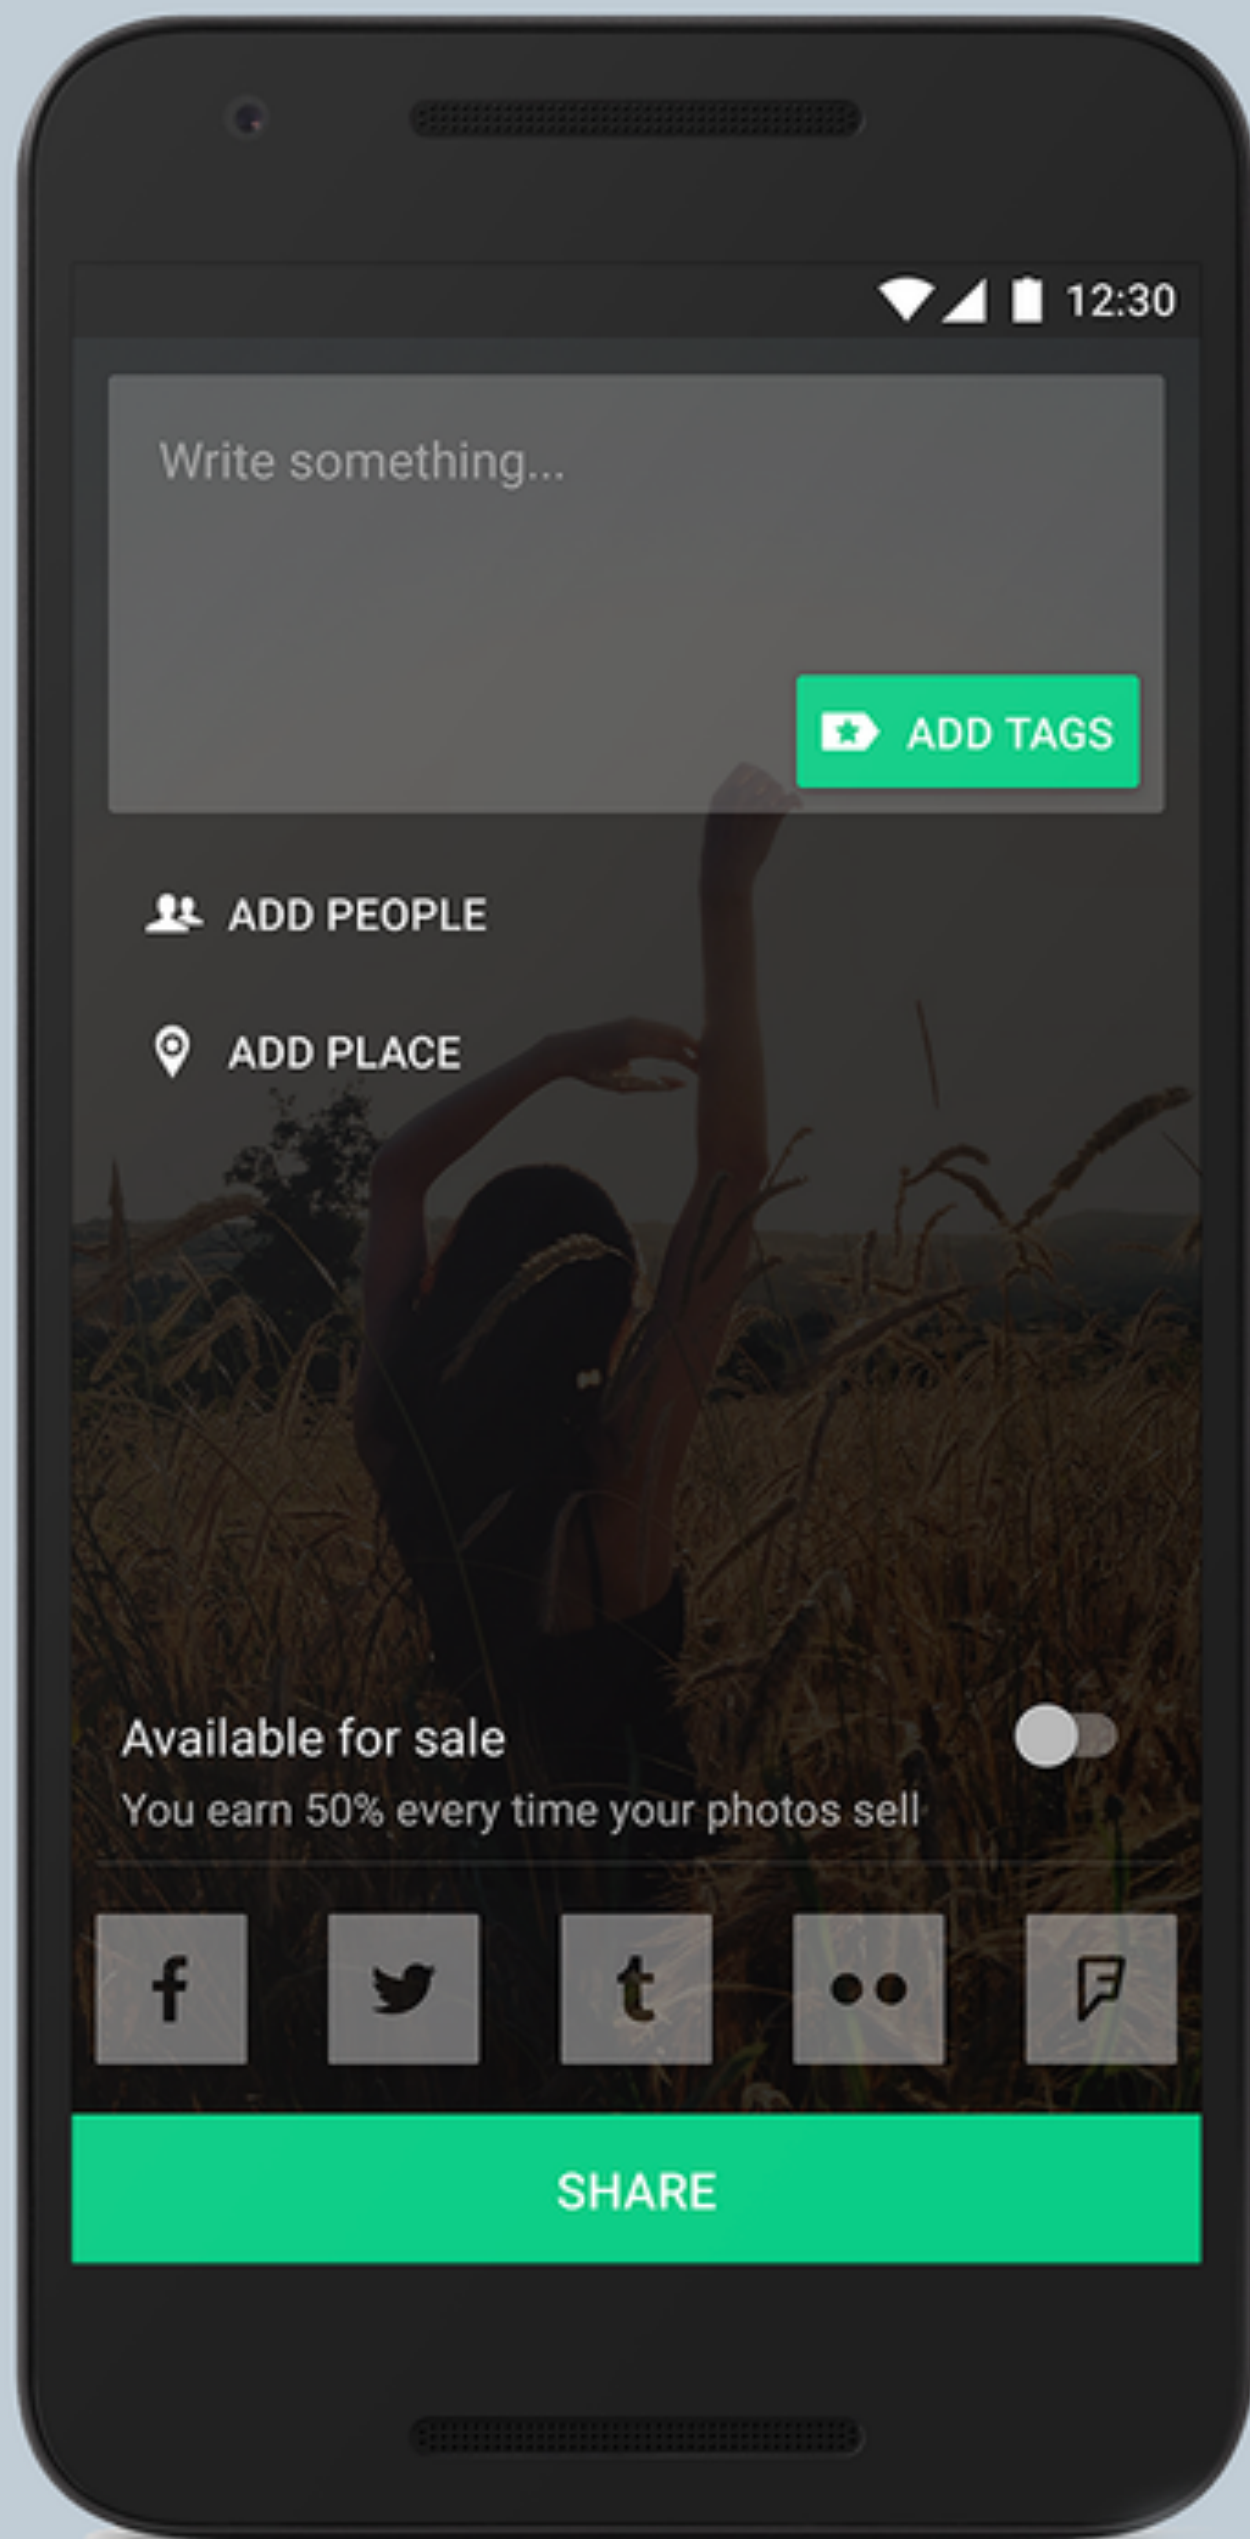



 Sari  245  47  

EyeEm

What we've automated Auto-tagging

A person with short, wavy pink hair and a bright red coat is holding a vintage-style camera with a flash up to their eye, as if taking a photo. They are standing on a sandy beach with the ocean in the background. The overall tone of the image is muted and artistic.



EyeEm

Problem:
Latency is too high

EyeEm

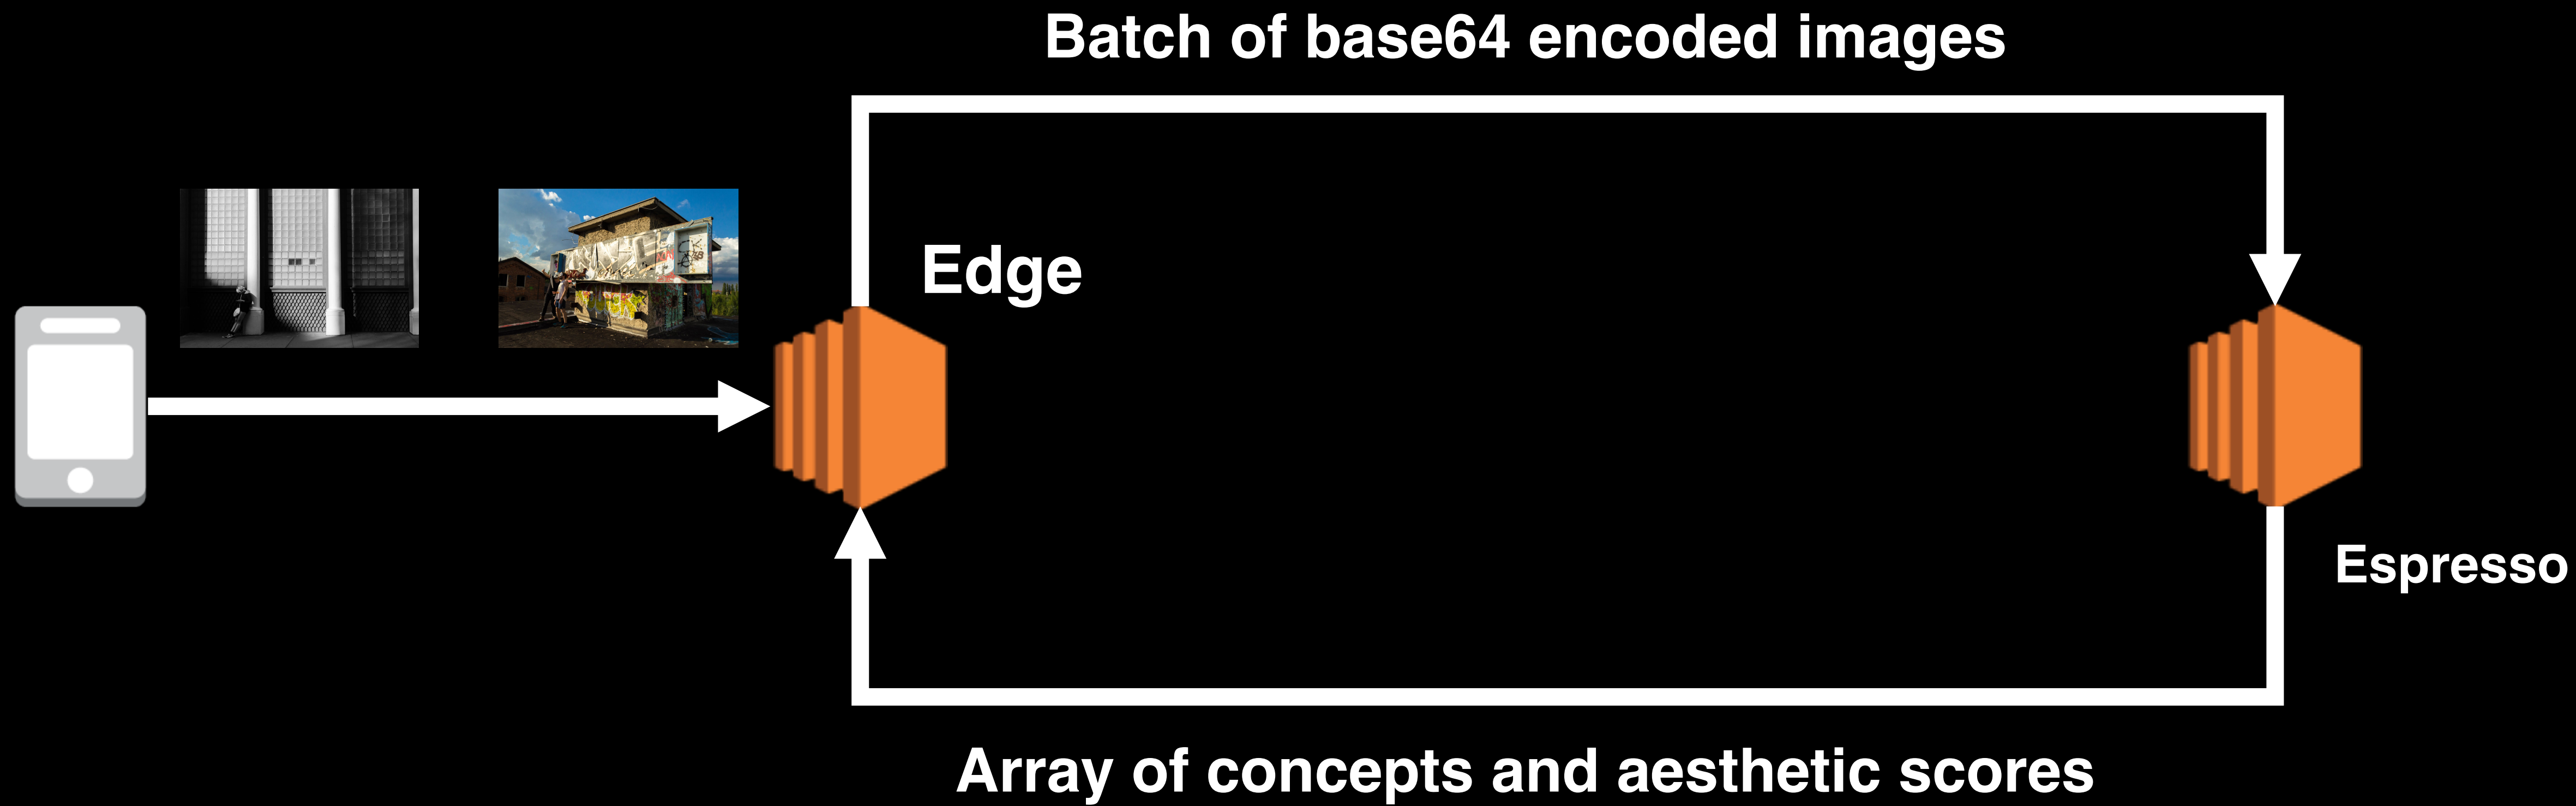
Problem:

This is getting expensive...

EyeEm

Batch inference

"GPU-Based Deep Learning Inference: A Performance and Power Analysis", NVIDIA
2015



EyeEm

Image reviewing pipeline

EyeEm

The traditional stock
photography
industry.



EyeEm



EyeEm

What we've automated

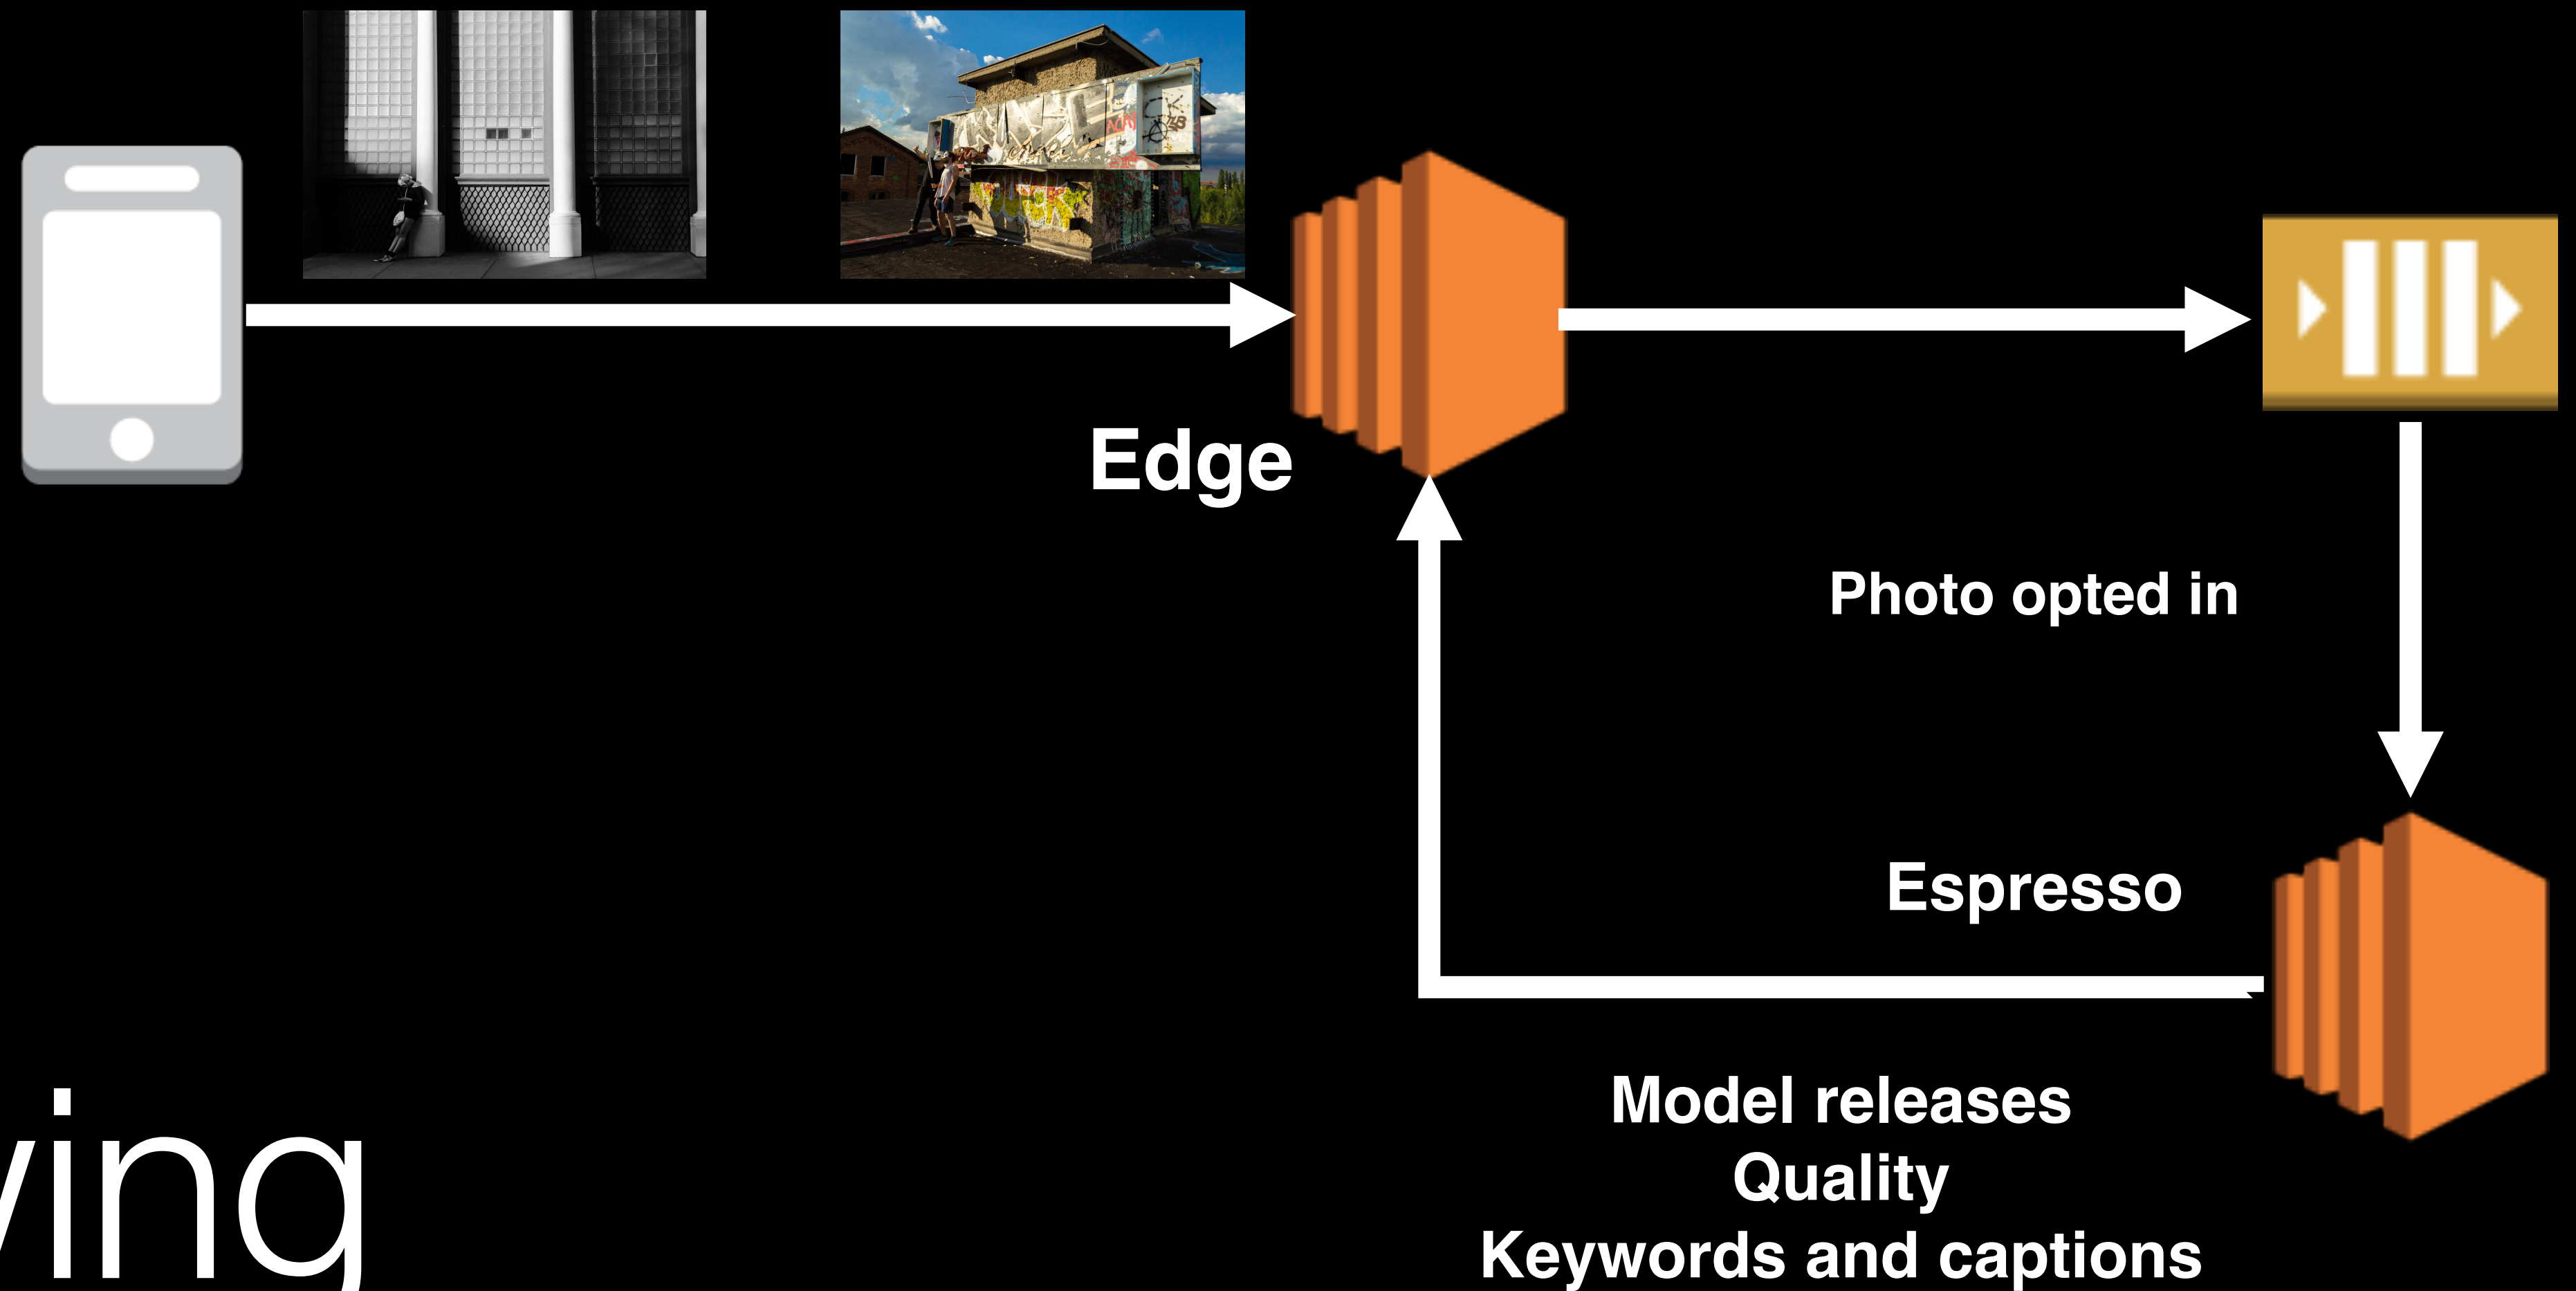
Keywording

Captioning

Photo quality rating

Model release detection

Photo reviewing process

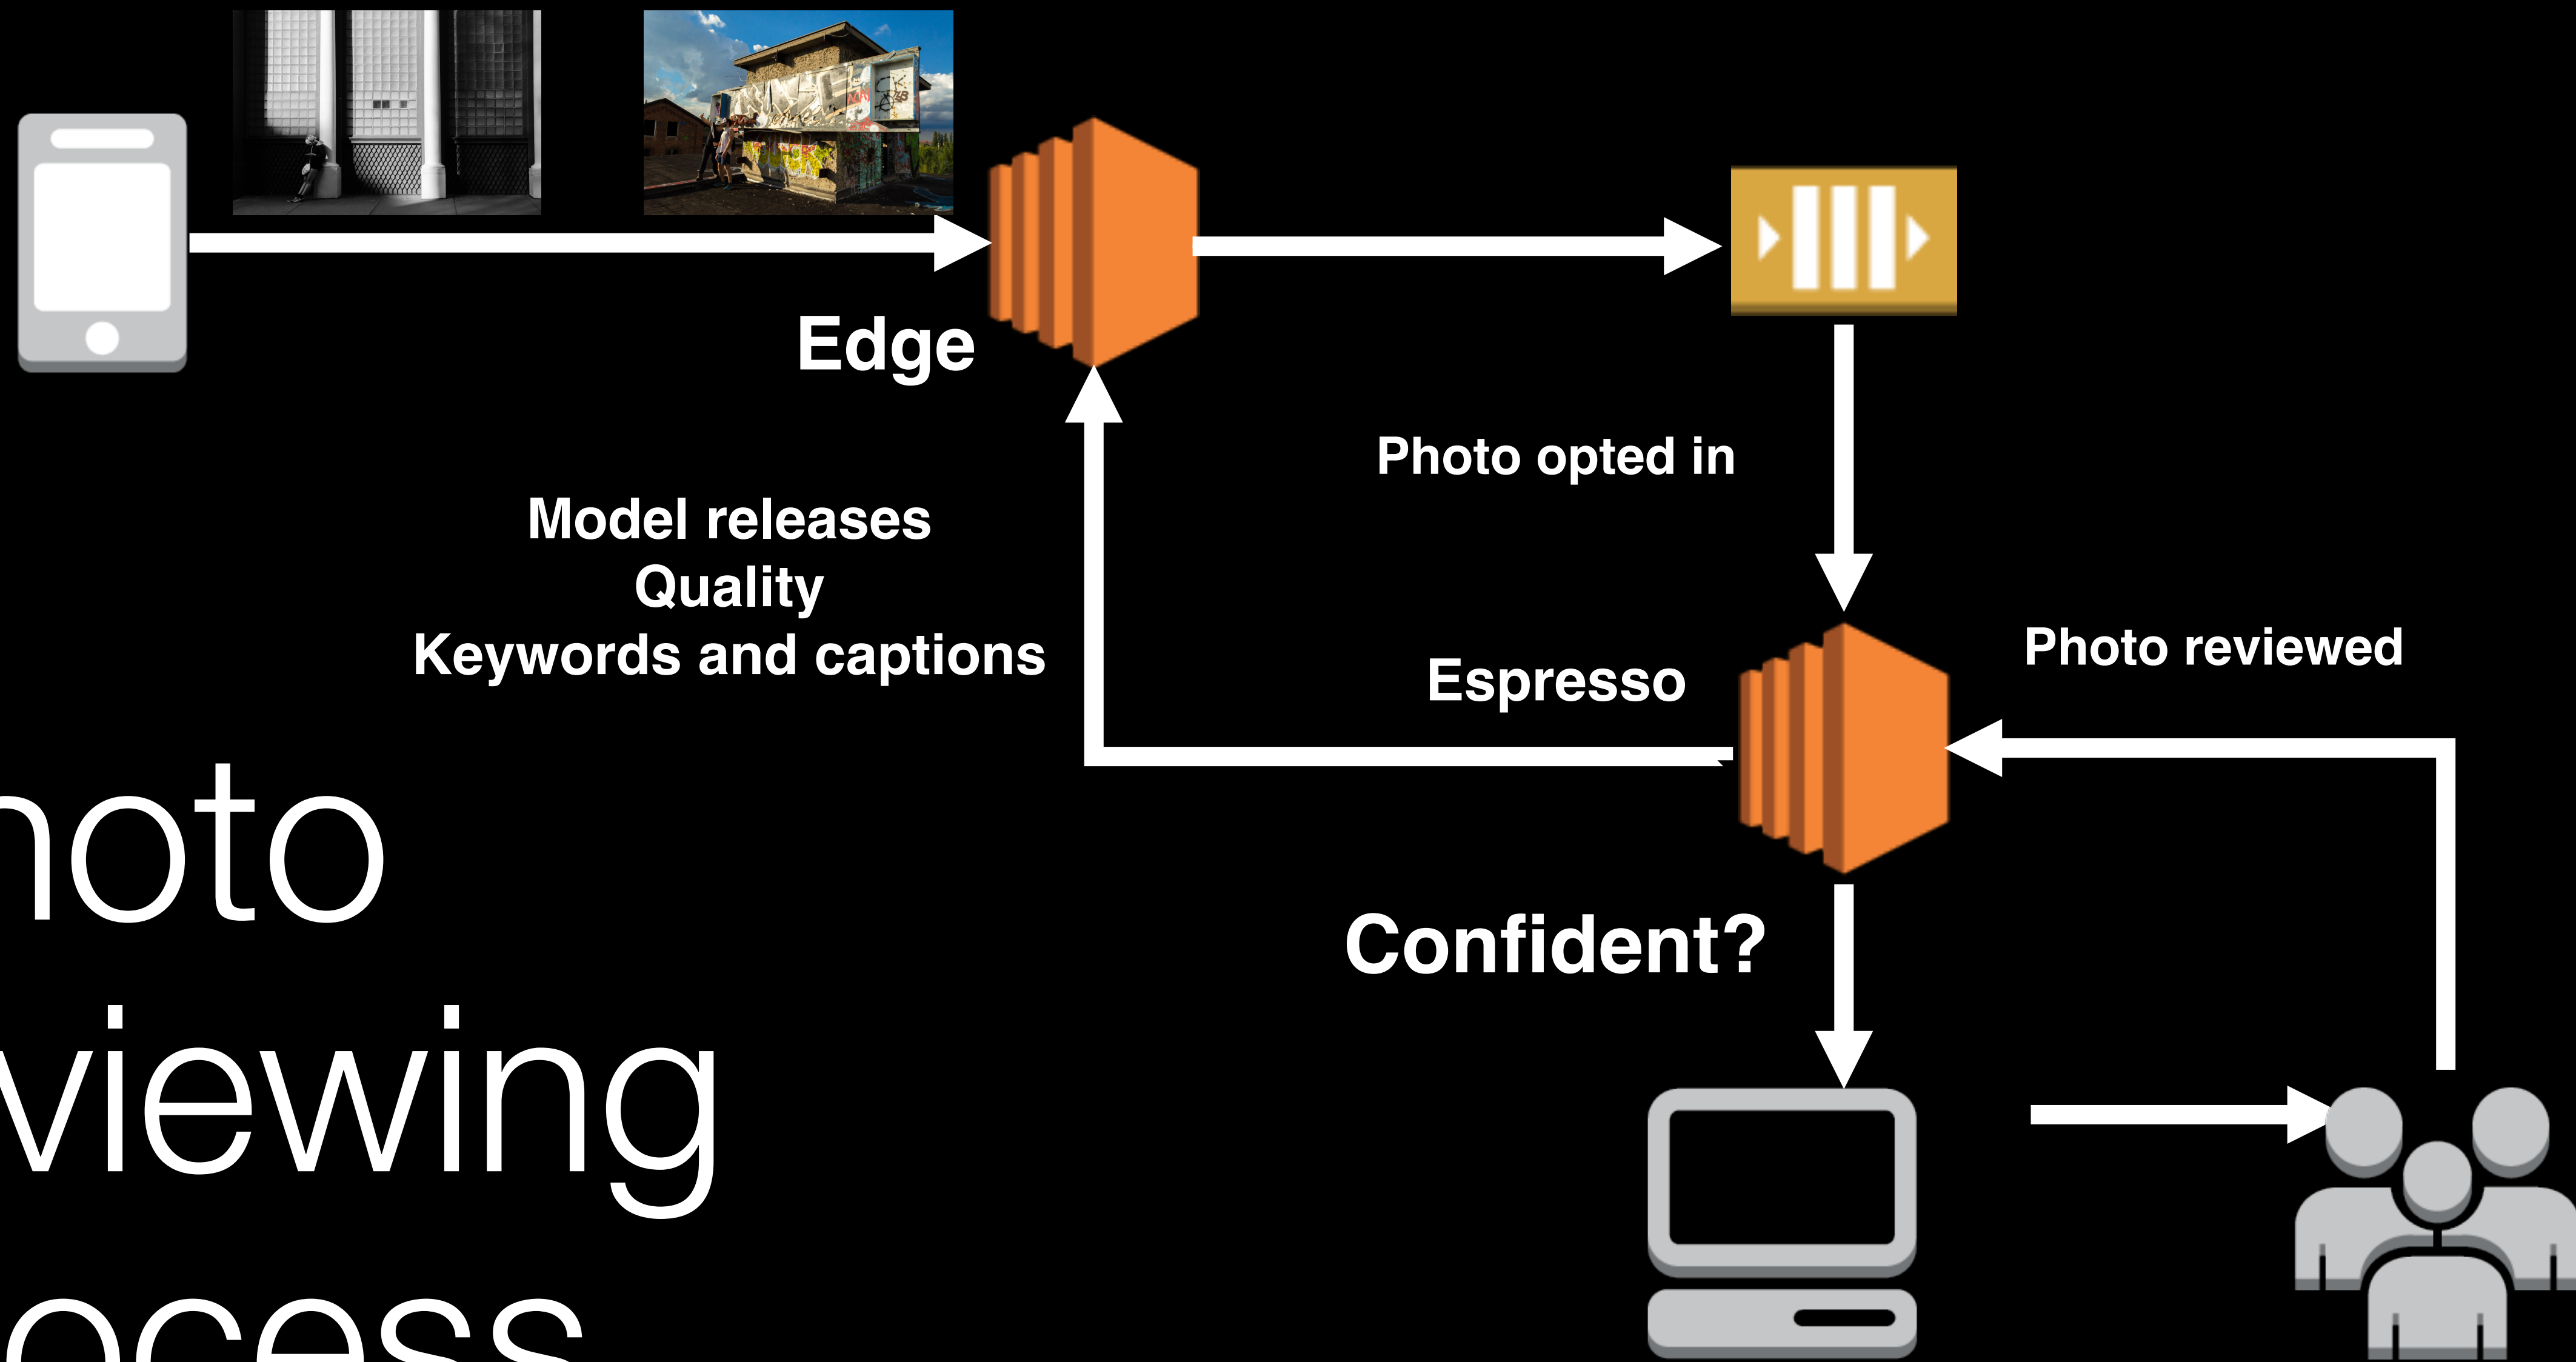






Woman photographing through smartphone

Photo reviewing process



A person with vibrant pink hair and a bright red coat is holding a vintage-style camera with a large external flash. They are standing on a beach with the ocean in the background. The image is semi-transparent, serving as a background for the text.

EyeEm

Asynchronous architectures are necessary to scale

Batch inference is essential for cost optimization and latency

90% accuracy can be misleading, but it can be optimized

EyeEm

Everything runs

EyeEm

End: Data

Data pipeline for analytics and machine learning



EyeEm

Espresso



aesthetics



concepts



quality



captions



**personalised
aesthetics**



...

EyeEm

Problem:
Too much data!

EyeEm

One photo

Keywords

Captions

Confidence

Quality

Model releases

Aesthetics

EyeEm

Concepts/Confidence

```
{ "photoId": "102394173", "concepts": [ [ "real people", 0.99792659282684326 ], [ "nature", 0.97948437929153442 ], [ "sky", 0.97900015115737915 ], [ "outdoors", 0.92300915718078613 ], [ "day", 0.92132925987243652 ], [ "beach", 0.91038638353347778 ], [ "scenics", 0.88209074735641479 ], [ "lifestyles", 0.84903037548065186 ], [ "two people", 0.82873058319091797 ], [ "leisure activity", 0.80806028842926025 ], [ "sand", 0.79079025983810425 ], [ "tranquility", 0.78834164142608643 ], [ "men", 0.78776693344116211 ], [ "full length", 0.78717267513275146 ], [ "beauty in nature", 0.69489455223083496 ], [ "landscape", 0.64603018760681152 ], [ "tree", 0.48800194263458252 ], [ "mammal", 0.31885015964508057 ] ], "confidence": 0.8376338593142145 }
```


EyeEm


Captions

```
{ "photoId": "93900842", "captions": [ { "caption": "STREAM FLOWING THROUGH  
ROCKS", "probability": 0.656147889422701, "wordsProbability": [ [ "STREAM", -  
0.47248415835501062 ], [ "FLOWING", -0.29332975536168115 ], [ "THROUGH", -  
0.59238967310011115 ], [ "ROCKS", -1.1257078392955155 ] ] }, { "caption": "VIEW OF  
STREAM FLOWING THROUGH  
ROCKS", "probability": 0.5243615522471853, "wordsProbability": [ [ "VIEW", -  
1.8661552402619594 ], [ "OF", -0.0041088119466486661 ], [ "STREAM", -  
0.087614953179970007 ], [ "FLOWING", -1.3746558803077191 ], [ "THROUGH", -  
0.31274626783455156 ], [ "ROCKS", -1.5081047112344315 ] ] }, { "caption": "VIEW OF  
STREAM IN  
FOREST", "probability": 0.4874233842518491, "wordsProbability": [ [ "VIEW", -  
1.8661552402619594 ], [ "OF", -0.0041088119466486661 ], [ "STREAM", -  
0.087614953179970007 ], [ "IN", -1.3630659191281458 ], [ "FOREST", -  
1.690672657028897 ] ] }, { "caption": "VIEW OF STREAM IN  
WATER", "probability": 0.43154049087804885, "wordsProbability": [ [ "VIEW", -  
1.8661552402619594 ], [ "OF", -0.0041088119466486661 ], [ "STREAM", -  
0.087614953179970007 ], [ "IN", -1.3630659191281458 ], [ "WATER", -  
2.3804022920975902 ] ] ] ] }
```


EyeEm

Billions of events

EyeEm



Apache Flume



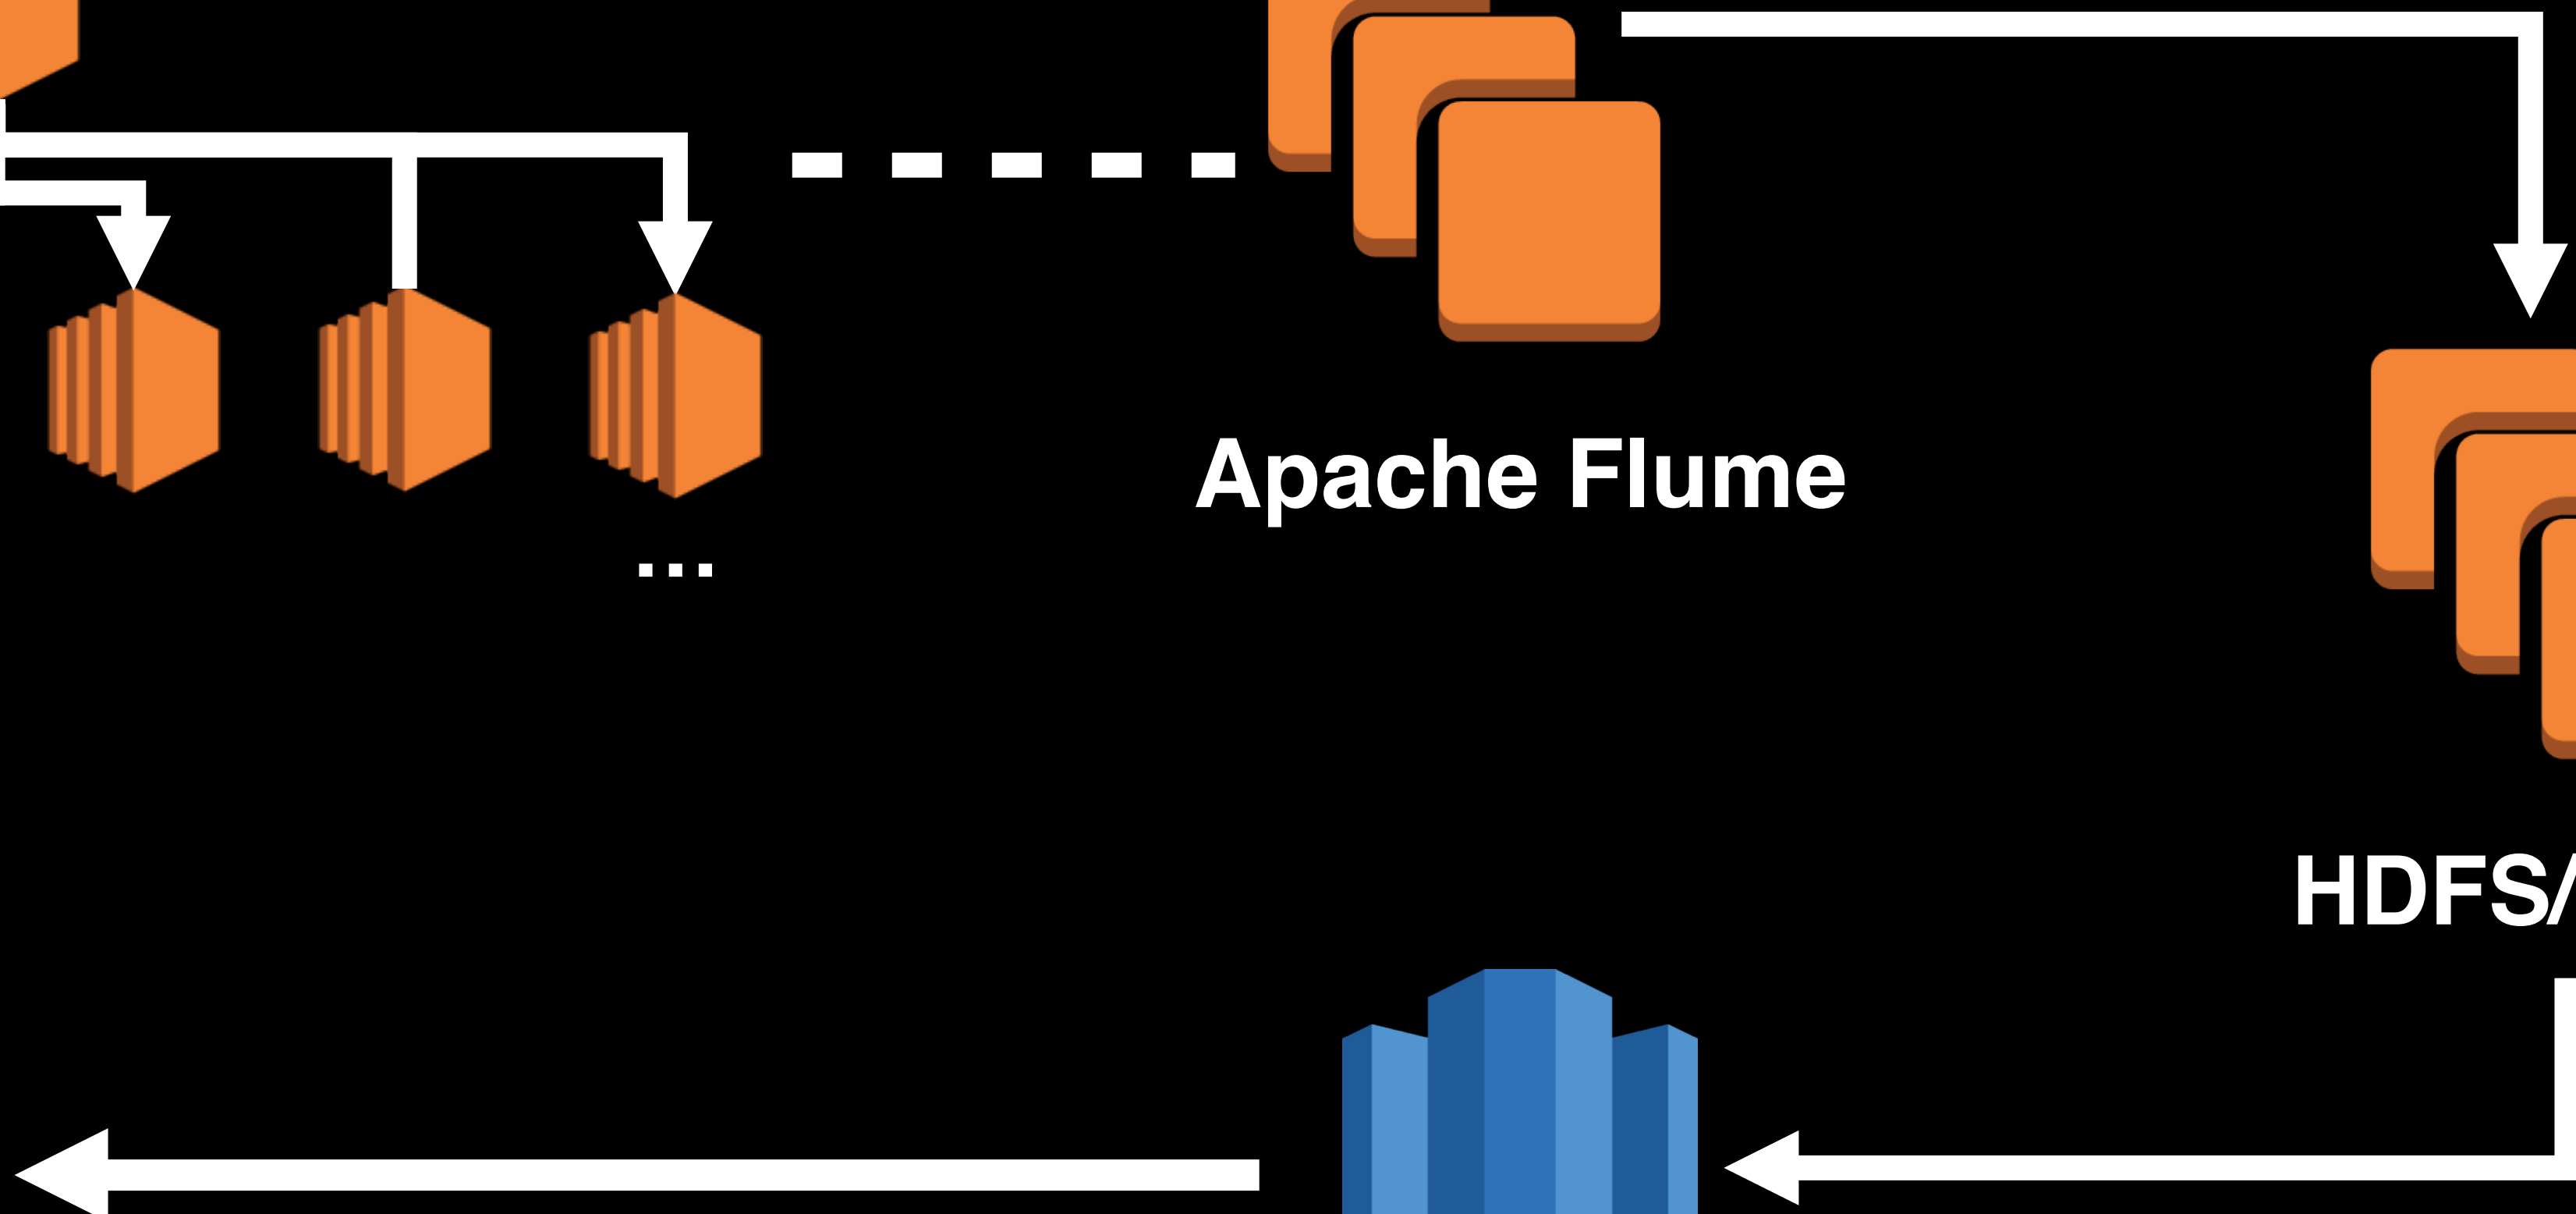
HDFS/Spark



Report



Redshift




EyeEm

Our move to EMR

EyeEm

Analytics pipeline:
Spark/EMR/Redshift

EyeEm



Apache Flume



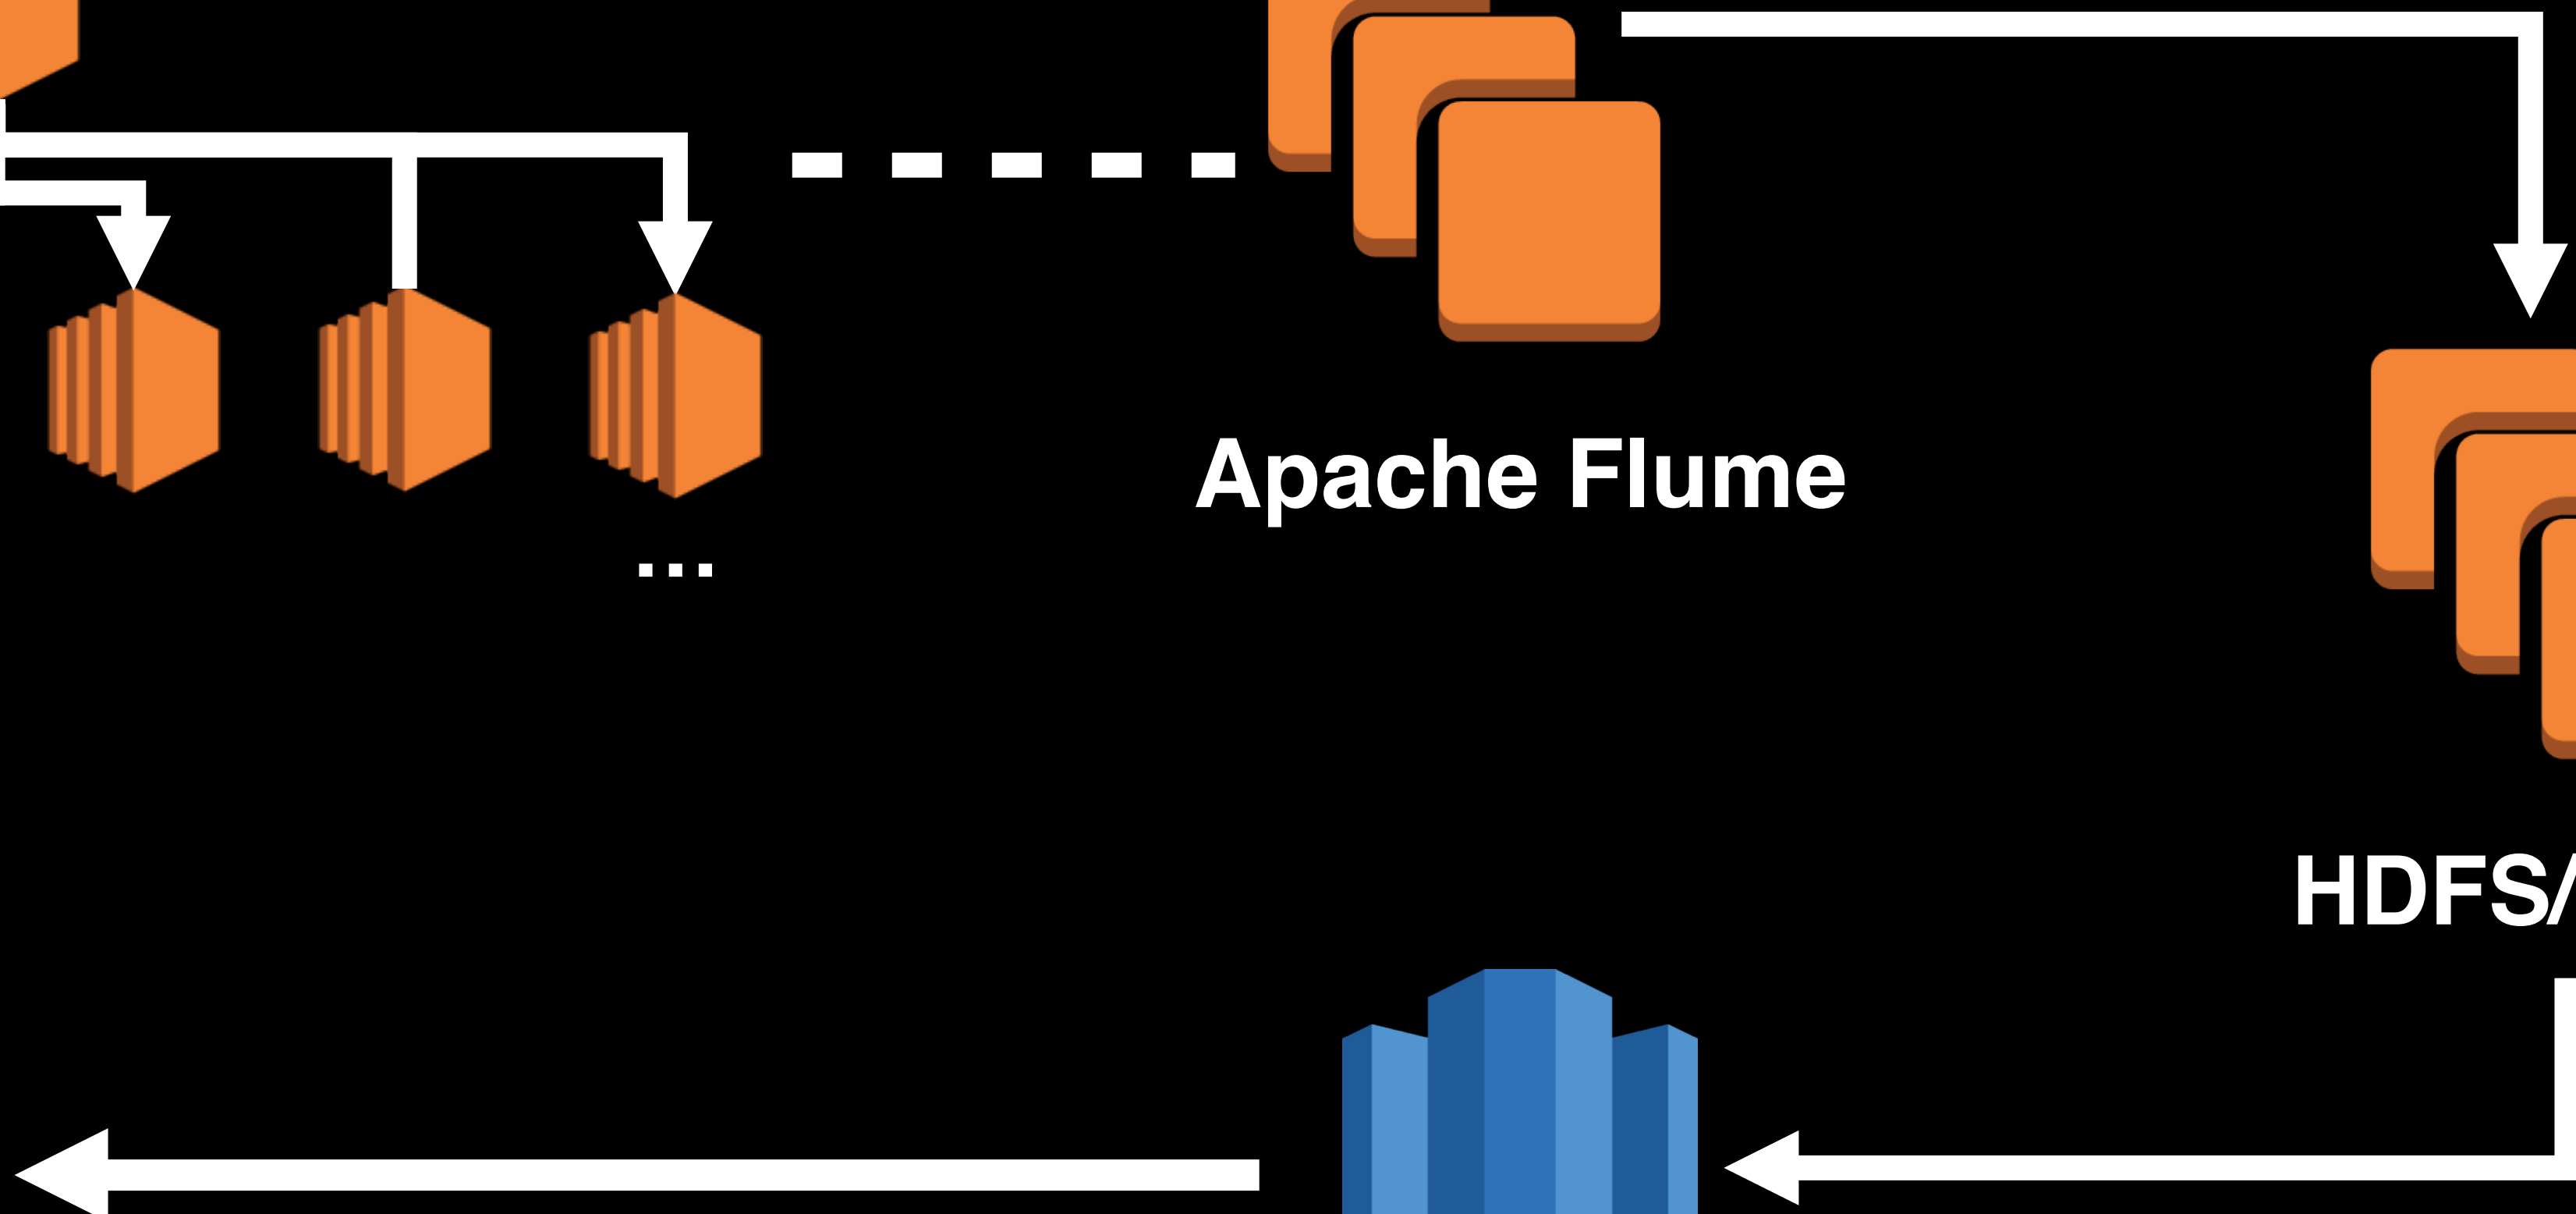
HDFS/Spark






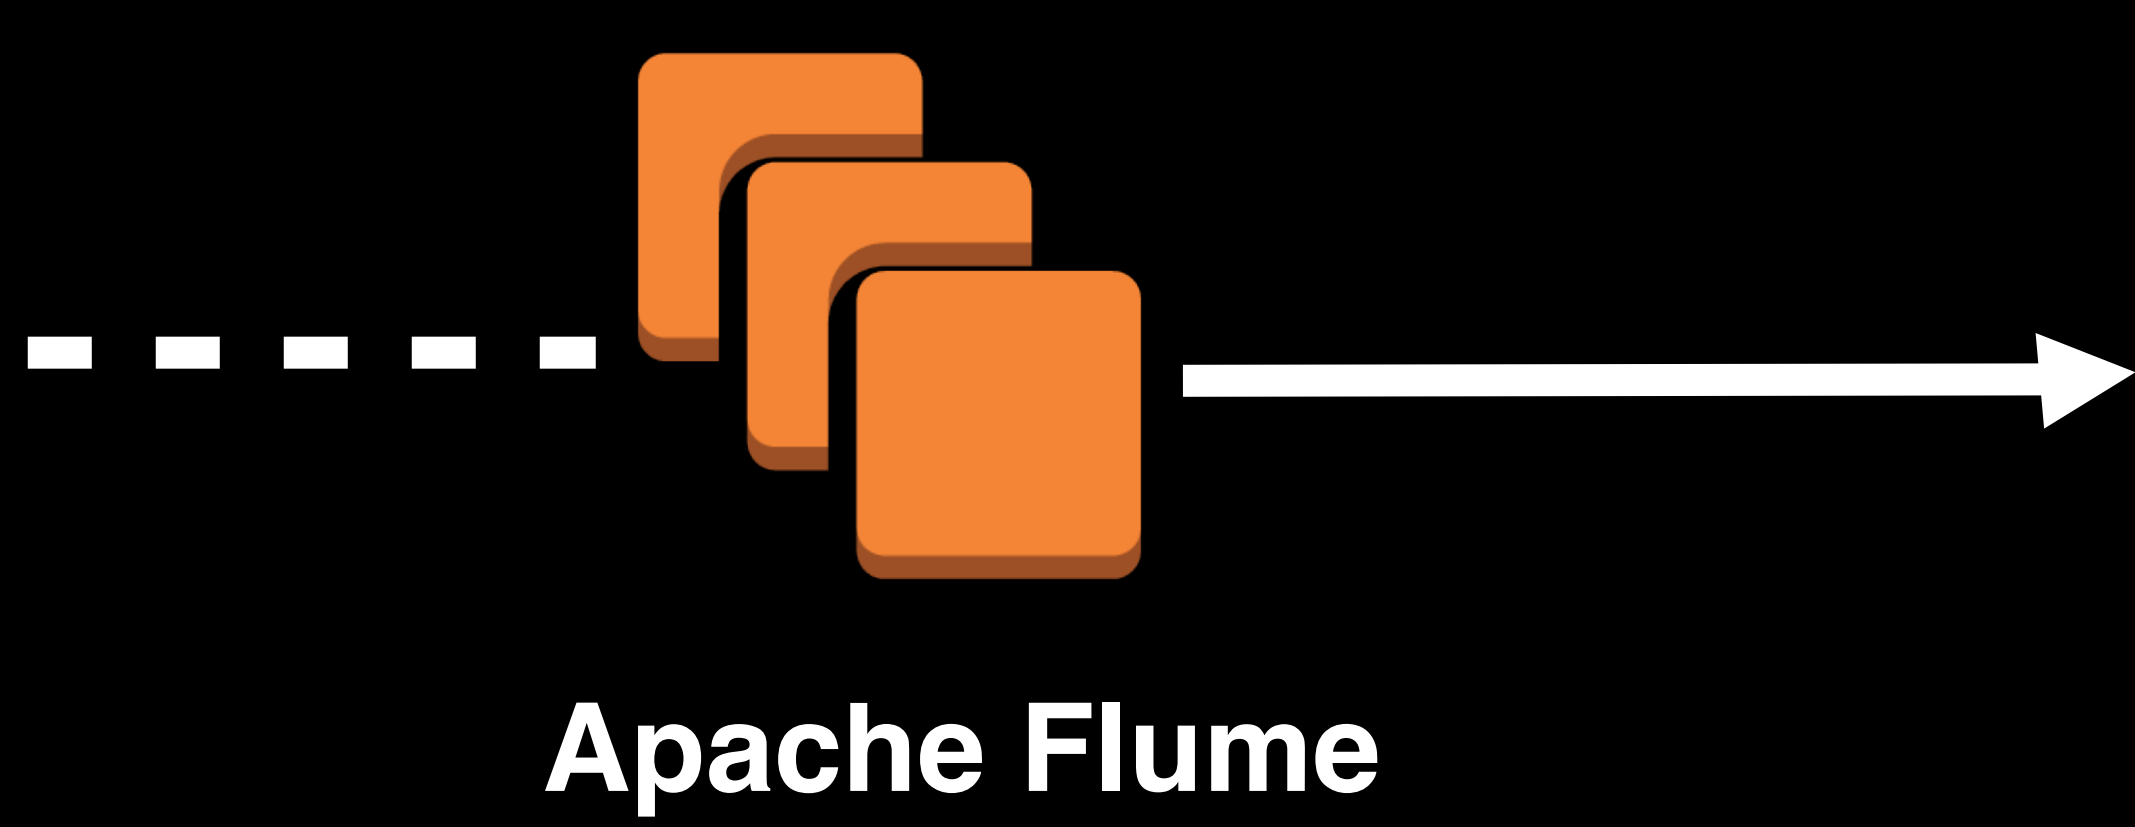
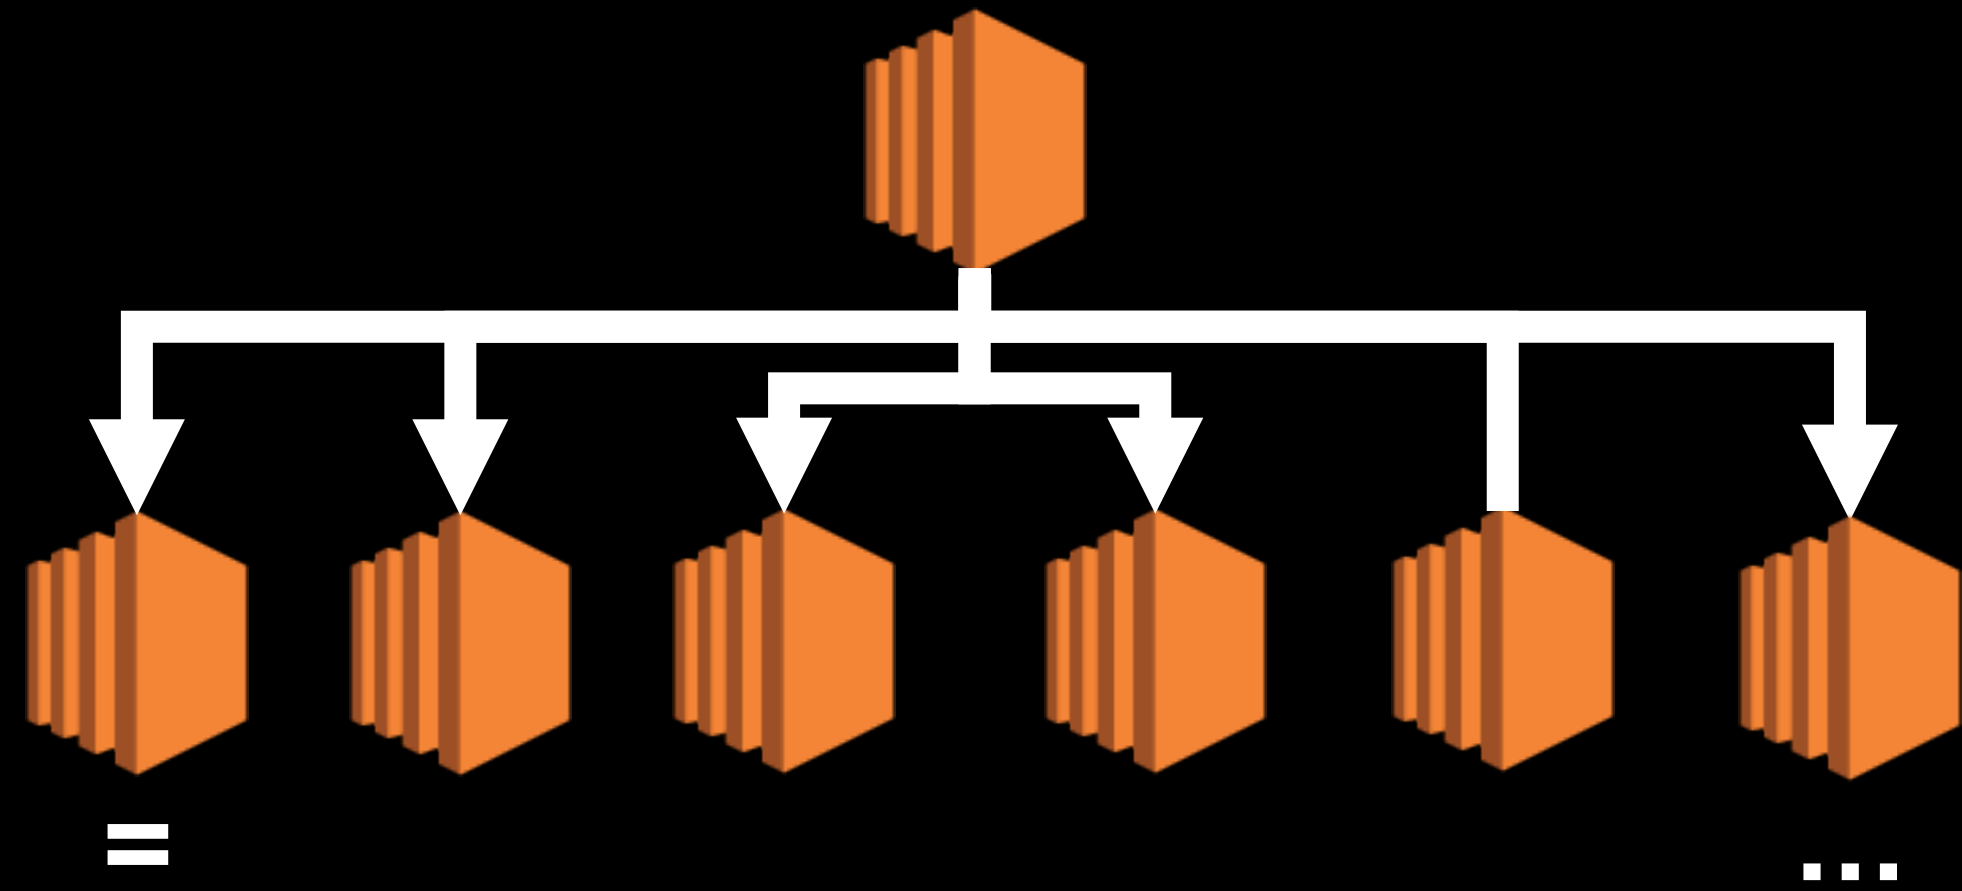
Report



Redshift



EyeEm



Report

Redshift


EMR



EyeEm

ML pipeline:
Spark ML/EMR/ElastiCache

EyeEm




Apache Flume

Events



EMR




Redshift




Report



EyeEm



Events




Apache Flume



Clients



API



ElastiCache



EMR



EyeEm

Machines produce a ton of data

Hadoop management is time consuming, avoid it!

Event driven data architectures have many advantages

EyeEm

Thank you!

@rinofm