


Safeguarding Video Content: AI-Powered Tokenization for Ironclad Media Delivery




Speaker introduction

About Me:



**David
Reisfeld**

Google Cloud

- Current - Manager Network Specialist  Google Cloud
- Past -
 - Product Management & Pre-Sales Engineering



Limelight LUMEN[®]





01

What's the problem?

Why do you care about video piracy?

**Digital video piracy
deprives the U.S. economy
of a minimum of \$29.2
billion in reduced revenue
each year***



* <https://www.uschamber.com/intellectual-property/quick-take-your-primer-digital-piracy-and-its-impact-the-us-economy>

Why worry about solving for video piracy?



**Loss of
Revenue**



**Cost of
Delivery**

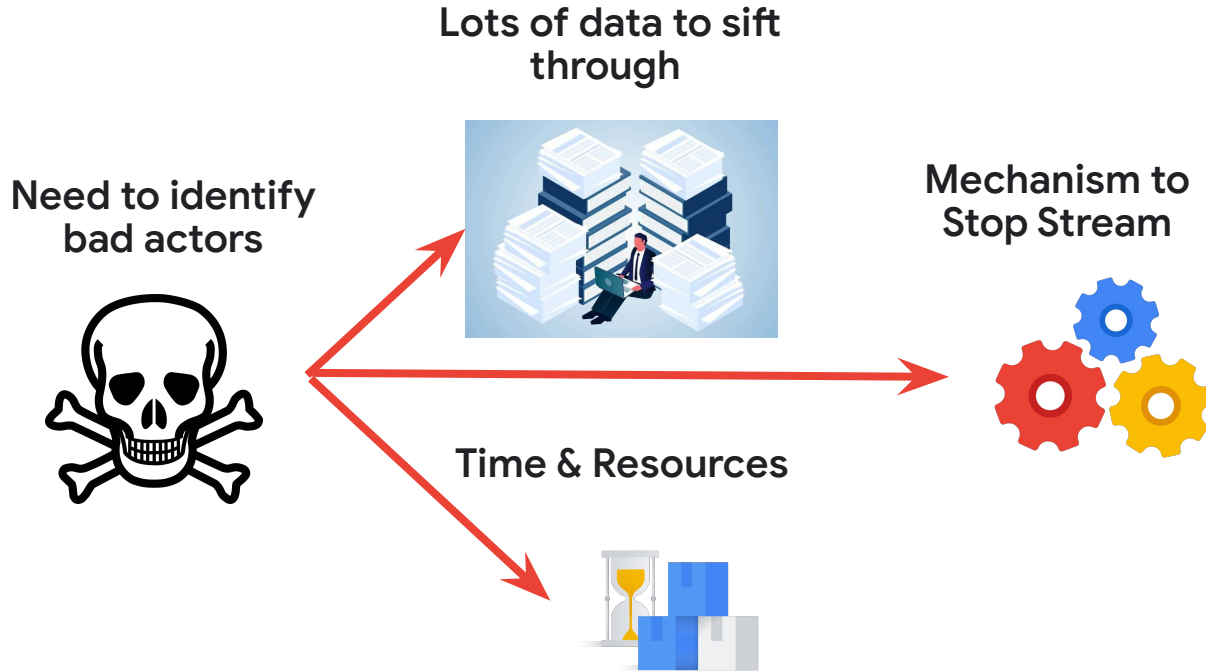



02

Just stop it

Block the bad actors

Just Stop It - Simple





How do Content Providers Try to Mitigate Piracy Today

Monitor known pirate websites





Manually Process Logs

Rotate access tokens





03

Lets automate the problem away

Use machine learning to help identify bad actors

Basic Workflow for Token Revocation




End
User



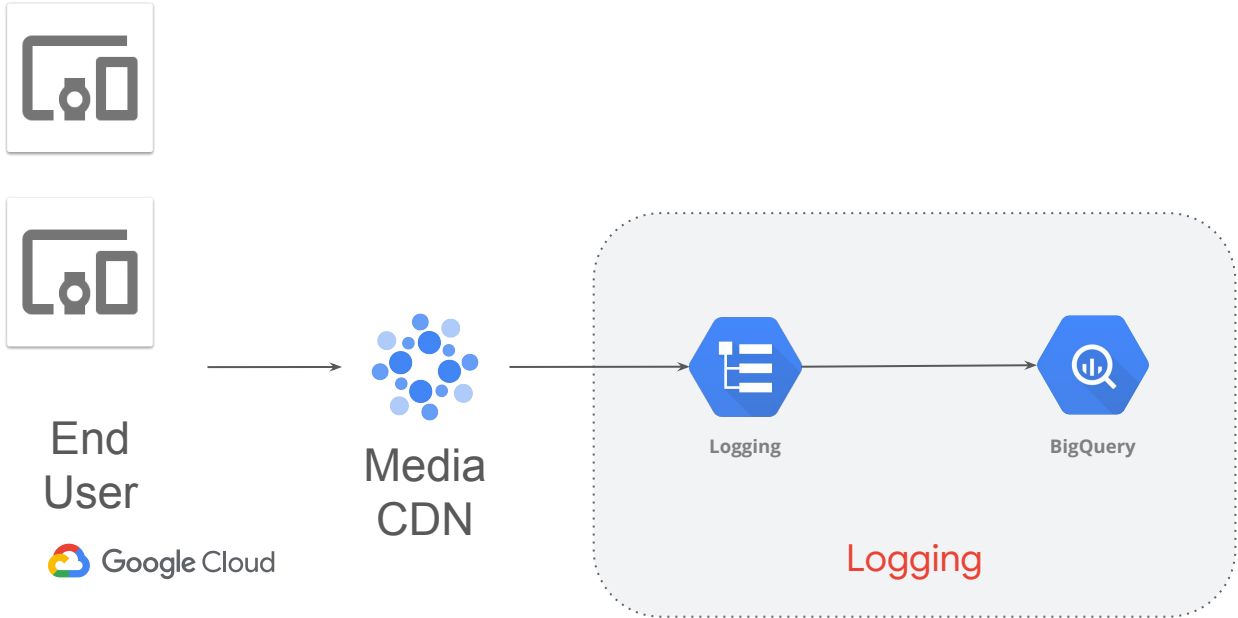
Token Revocation Workflow



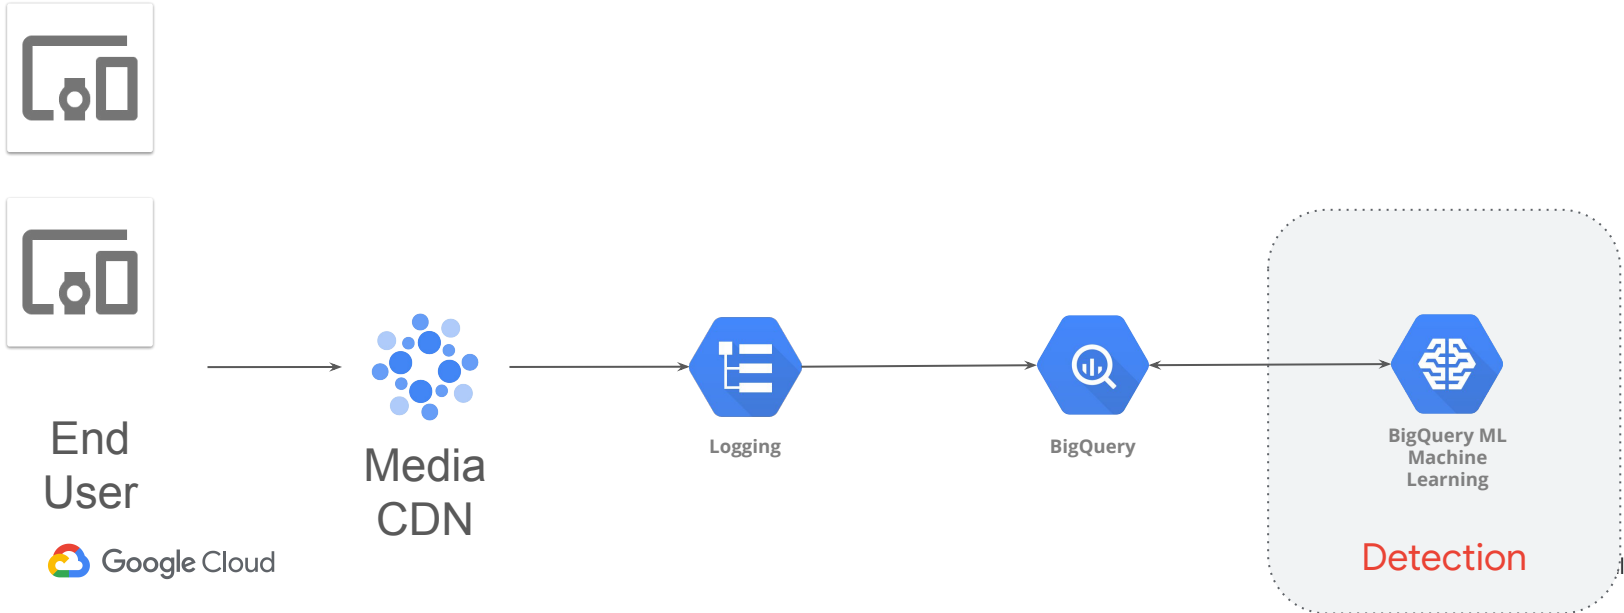
End
User



Token Revocation Workflow



Token Revocation Workflow

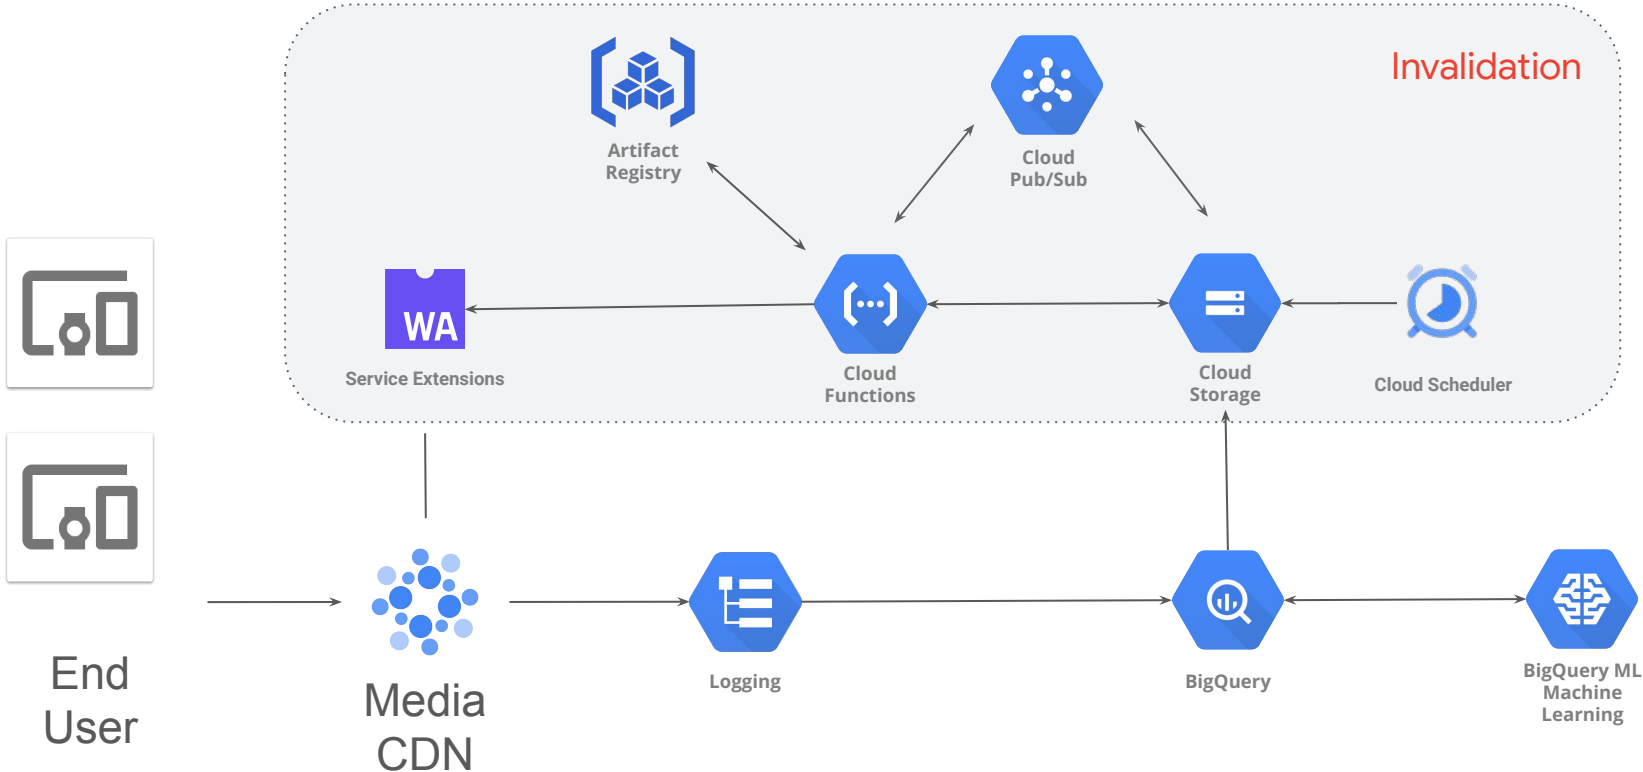


Leverage details from the Media CDN log lines

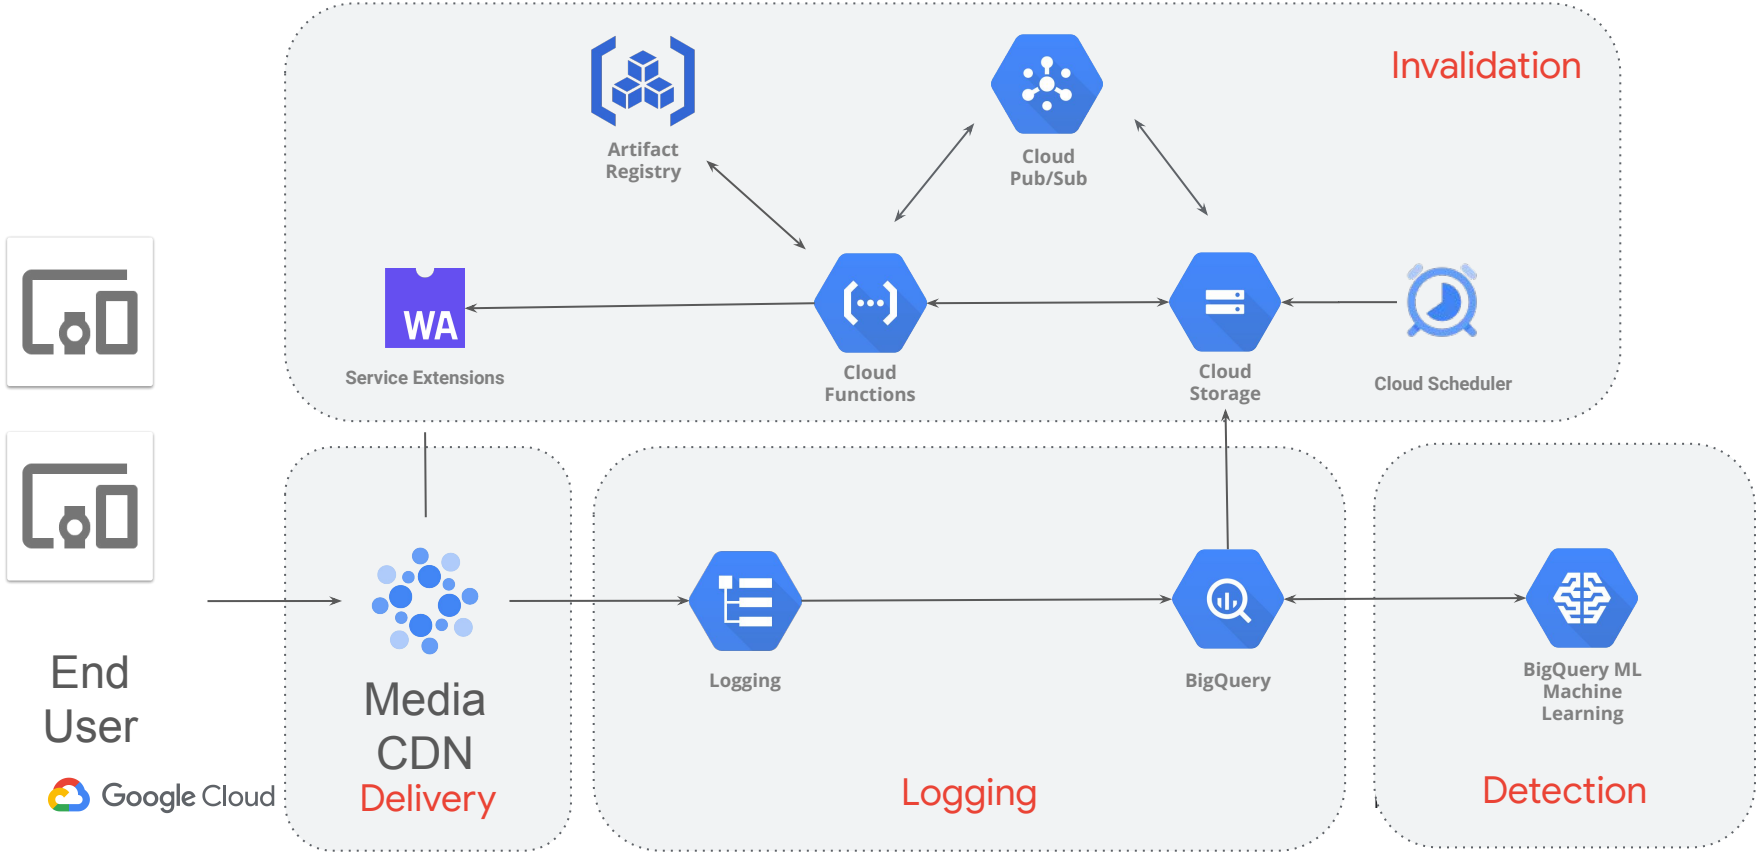
```
agent: "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/120.0.6322.59 Safari/537.36"
insertId: "695f8dd6-0000-2fa4-9f8a-4a1c1c1c1c1c"
jsonPayload: {
  @type: "type.googleapis.com/google.cloud.edgecache.v1.EdgeCacheLogEntry"
  cacheId: "ord,ord-f951c4f4, lga"
  cacheKeyFingerprint: "f5b290caa70326ee"
  cacheMode: "FORCE_CACHE_ALL"
  cacheStatus: "stale,miss,uncacheable"
  clientAsn: "701"
  clientCity: "Bedminster"
  clientRegionCode: "US"
  enforcedSecurityPolicy: {
    configuredAction: "ACCEPT"
    name: "block-bad-countries"
    outcome: "ACCEPT"
    priority: 999
  }
  httpTtfb: "0.089000201s"
  latency: "0.089038891s"
  metroIataCode: "LGA"
  origin: "lnmp-7"
  originIp: "34.149.39.34"
  proxyRegionCode: "US"
}
```

```
httpRequest: {
  protocol: "HTTP/2"
  remoteIp: "2600:802:c00::a"
  requestMethod: "GET"
  requestSize: "4428"
  requestUrl: "https://[redacted]images/bailey.jpg"
  responseSize: "407"
  status: 304
  userAgent: "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/120.0.6322.59 Safari/537.36"
}
insertId: "695f8dd6-0000-2fa4-9f8a-883d24f9e3b0a1"
jsonPayload: {
  @type: "type.googleapis.com/google.cloud.edgecache.v1.EdgeCacheLogEntry"
  cacheId: "ord,ord-f951c4f4, lga"
  cacheKeyFingerprint: "f5b290caa70326ee"
  cacheMode: "FORCE_CACHE_ALL"
  cacheStatus: "stale,miss,uncacheable"
  clientAsn: "701"
  clientCity: "Bedminster"
  clientRegionCode: "US"
  enforcedSecurityPolicy: {
    configuredAction: "ACCEPT"
    name: "block-bad-countries"
    outcome: "ACCEPT"
    priority: 999
  }
  httpTtfb: "0.089000201s"
  latency: "0.089038891s"
  metroIataCode: "LGA"
  origin: "lnmp-7"
  originIp: "34.149.39.34"
  proxyRegionCode: "US"
  proxyStatus: "Google-Edge-Cache"
  requestId: "fd980e3a-8ce1-4ad8-bc89-5d919b2065ac"
  tlsVersion: "TLS 1.3"
}
logName: "projects/[redacted]/logs/edgecache.googleapis.com%2Fedge_cache_request"
receiveTimestamp: "2024-04-13T21:23:19.439065659Z"
resource: {
  labels: {
    location: "us-west1"
    matched_path: "/images/*"
    path_matcher_name: "routes"
    resource_container: "projects/[redacted]"
    route_destination: "projects/[redacted]/locations/global/edgeCacheOrigins/lnmp-7"
    route_type: "ORIGIN"
    service_name: "[redacted]"
  }
}
```

Token Revocation Workflow



Token Revocation Workflow



RESULT:

No Human in the Loop





04 What can go wrong?

Some pitfalls with using automated invalidation

What can go wrong?

Time Interval


- Duration time of playback token
- Live event vs VOD
- Pipeline can only run as fast as the weakest link



What can go wrong?

Not looking at the right metrics

- Just looking at tokens
- Not looking at Geo or ASN
- Putting too much emphasis on ASN




What can go wrong?

Cost

- Log processing
 - Multi-CDN log ingestion
- Log storage
- Cost to run the invocation pipeline



Q&A



Thank you

Google Cloud

