SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

# How Video Analytics is Changing the Way We Store Video

Live Webcast

December 2, 2020

10:00 am PT

# Today's Presenters

**Jim Fister**
**Principal**
**The Decision Place**

**Glyn Bowden**
**CTO, AI & Data Practice**
**HPE**

**Kevin Cone**
**Media Analytics Segment Manager**
**Intel**

 

# SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

 

# SNIA-At-A-Glance

**185**
industry leading organizations

**2,000**
active contributing members

**50,000**
IT end users & storage pros worldwide

 

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

# What We Do

**Educate** vendors and users on cloud storage, data services and orchestration

**Support & promote** business models and architectures: OpenStack, Software Defined Storage, Kubernetes, Object Storage

**Understand** Hyperscaler requirements Incorporate them into standards and programs

**Collaborate** with other industry associations

# Agenda

- Ye olde video platforms
- The benefits of video analytics and the new use of video
- Video platform eco-system
- Security & Governance with analytics
- Summary

 

# Ye Olde Video Platforms

## Yes, they still exist. A lot!

 

# Ye Olde Video Platform

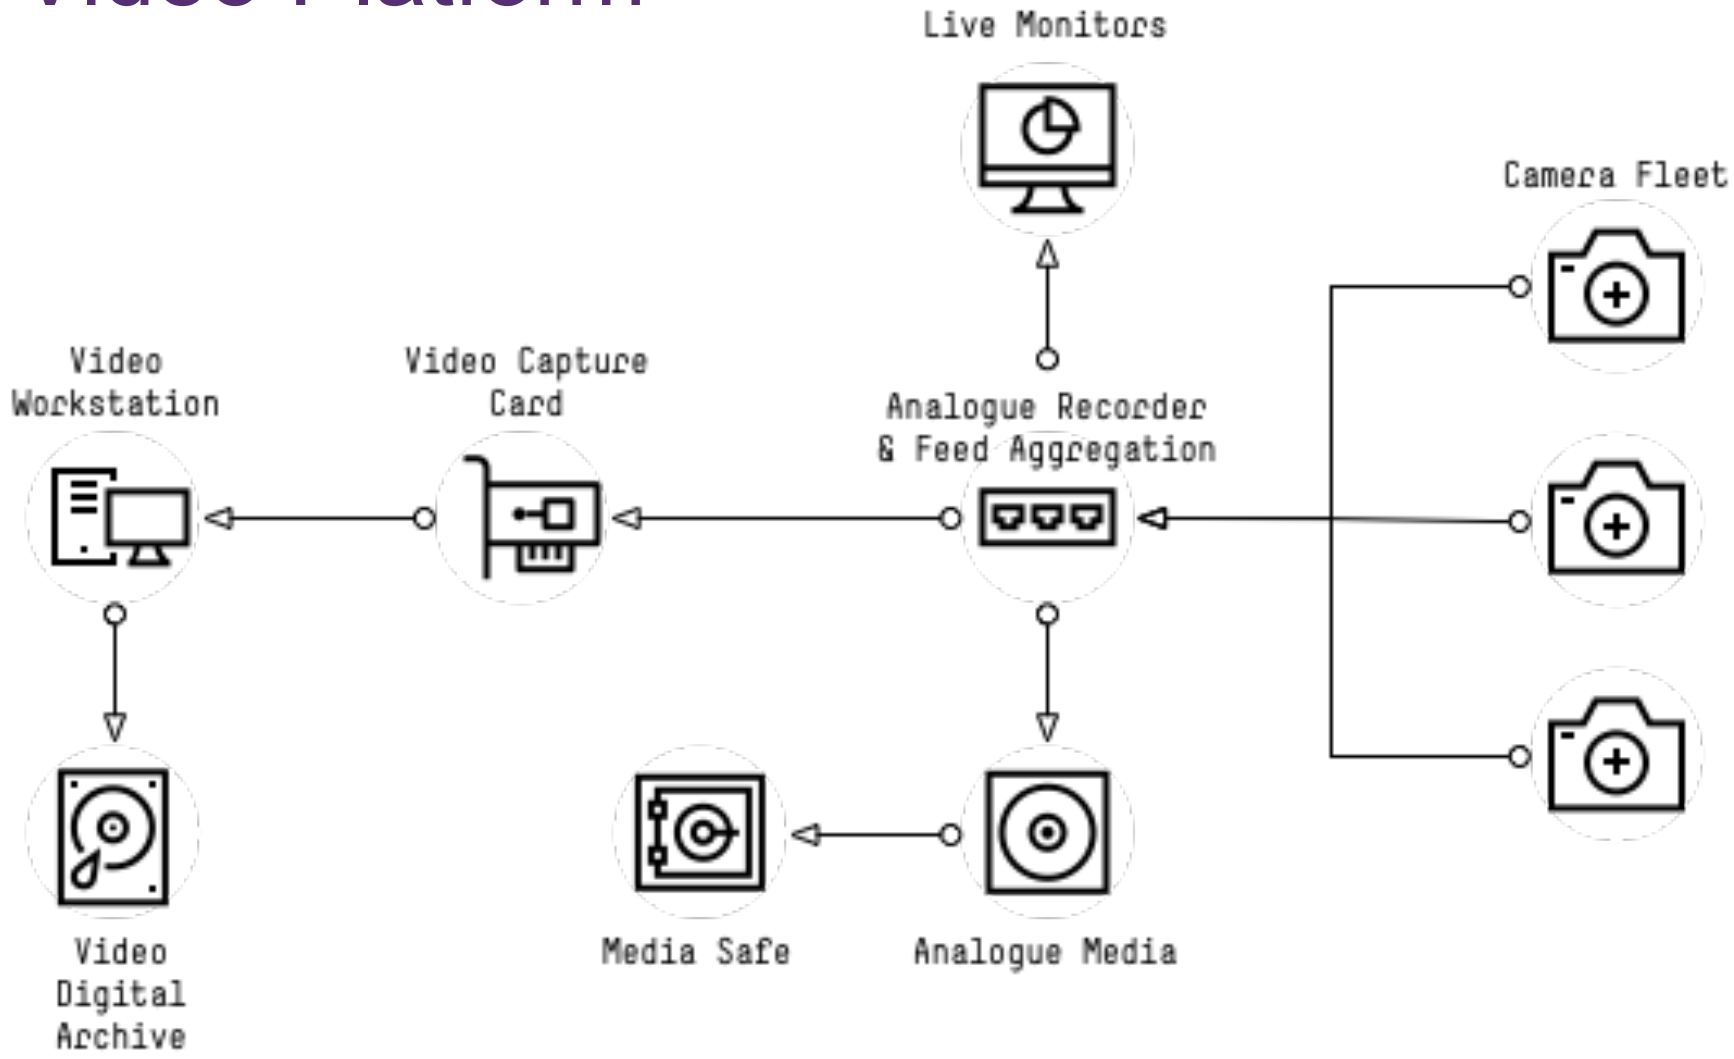


 

# Ye Olde Video Platform

- Analogue
- Stored in very low resolution
- Noisy / Grainy frames
- Often monochrome
- Speed of medium to ingest is Slooooow
  - When the media is even onsite!
- No indexing (other than little cards in boxes)

 

# The benefits of video analytics and the new use of video

## Fast Forward to today's use of video

 

# Open Visual Cloud Project

Extensible Software Architecture: Accelerating Services Innovation

**REFERENCE PIPELINES**

CDN Transcode | Intelligent Ad-Insertion | Video Conferencing | Interactive raytracing | Smart city | Library Curation | Game streaming | 360 streaming


**INDUSTRY FRAMEWORKS**

FFMPEG, GSTREAMER, NGINX

**OPTIMIZED INGREDIENTS**

Scalable video technology, x264, x265 | OpenVINO™ toolkit | Intel® OneAPI rendering toolkit | Scalable video technology

**CORE FUNCTIONS**

Decode | Inference | Render | Encode

OneAPI

Seamless support for Intel's Hardware and Software portfolio for visual cloud

GitHub

docker

Scale Media Workloads from Cloud to Edge

Github.com/OpenVisualCloud

Optimization Notice

 

# Smart City Example Pipeline

Demo at: https://youtu.be/BWU0SEqEfbo

Download from: https://github.com/OpenVisualCloud/Smart-City-Sample

- Effective detection of pedestrian, automobiles, and bicycles for street flow management
- Automated demographics: public space / transportation usage surveys to create smart cities
- Safety and Surveillance: traffic control and monitoring
- Fast detection allows for low latency and real time analytics in an edge environment.
- Includes smart upload for further analysis and archival of critical data

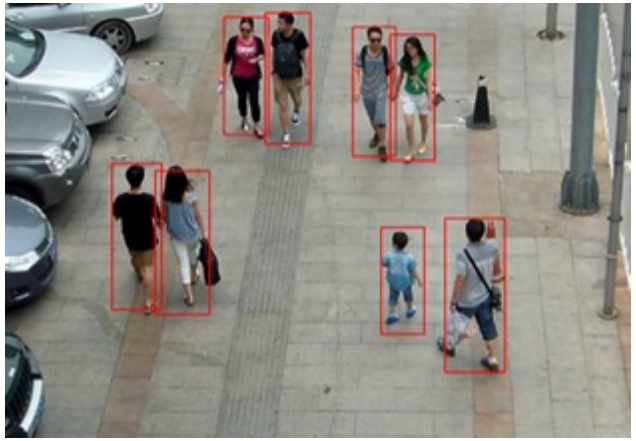

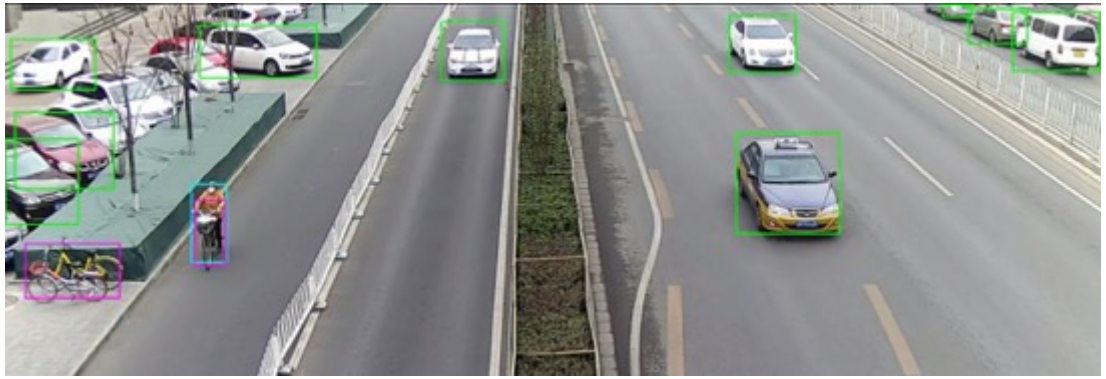

Sources: (TOP) https://github.com/opencv/open_model_zoo/blob/master/intel_models/person-vehicle-bike-detection-crossroad-0078/description/person-vehicle-bike-detection-crossroad-0078.md, (BOTTOM) https://github.com/opencv/open_model_zoo/blob/master/intel_models/person-detection-retail-0013/description/person-detection-retail-0013.md

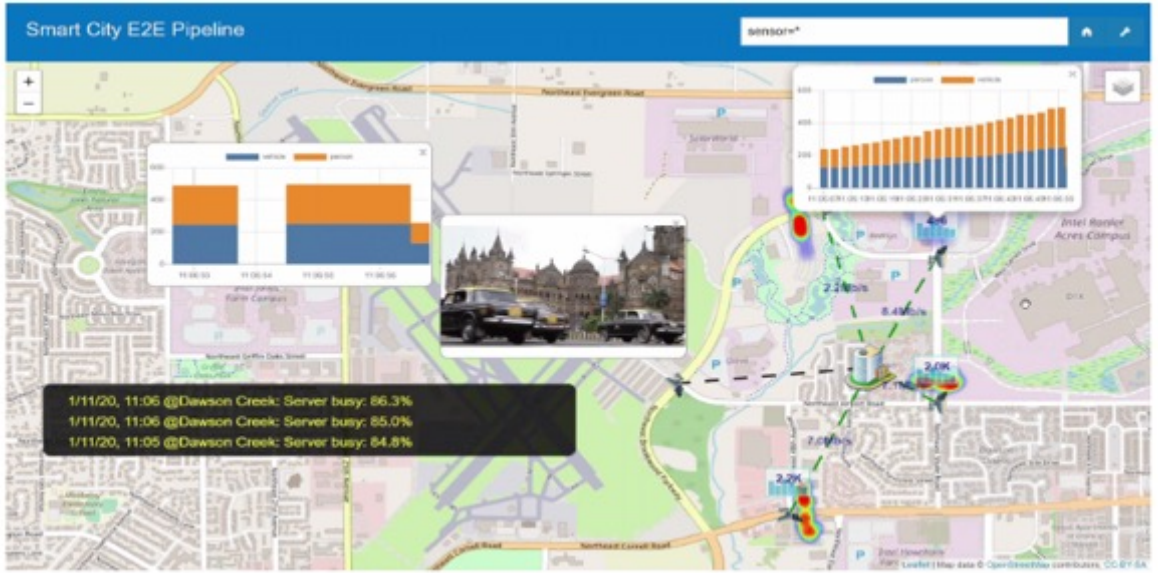


 

# Large Scale Video Curation System Overview

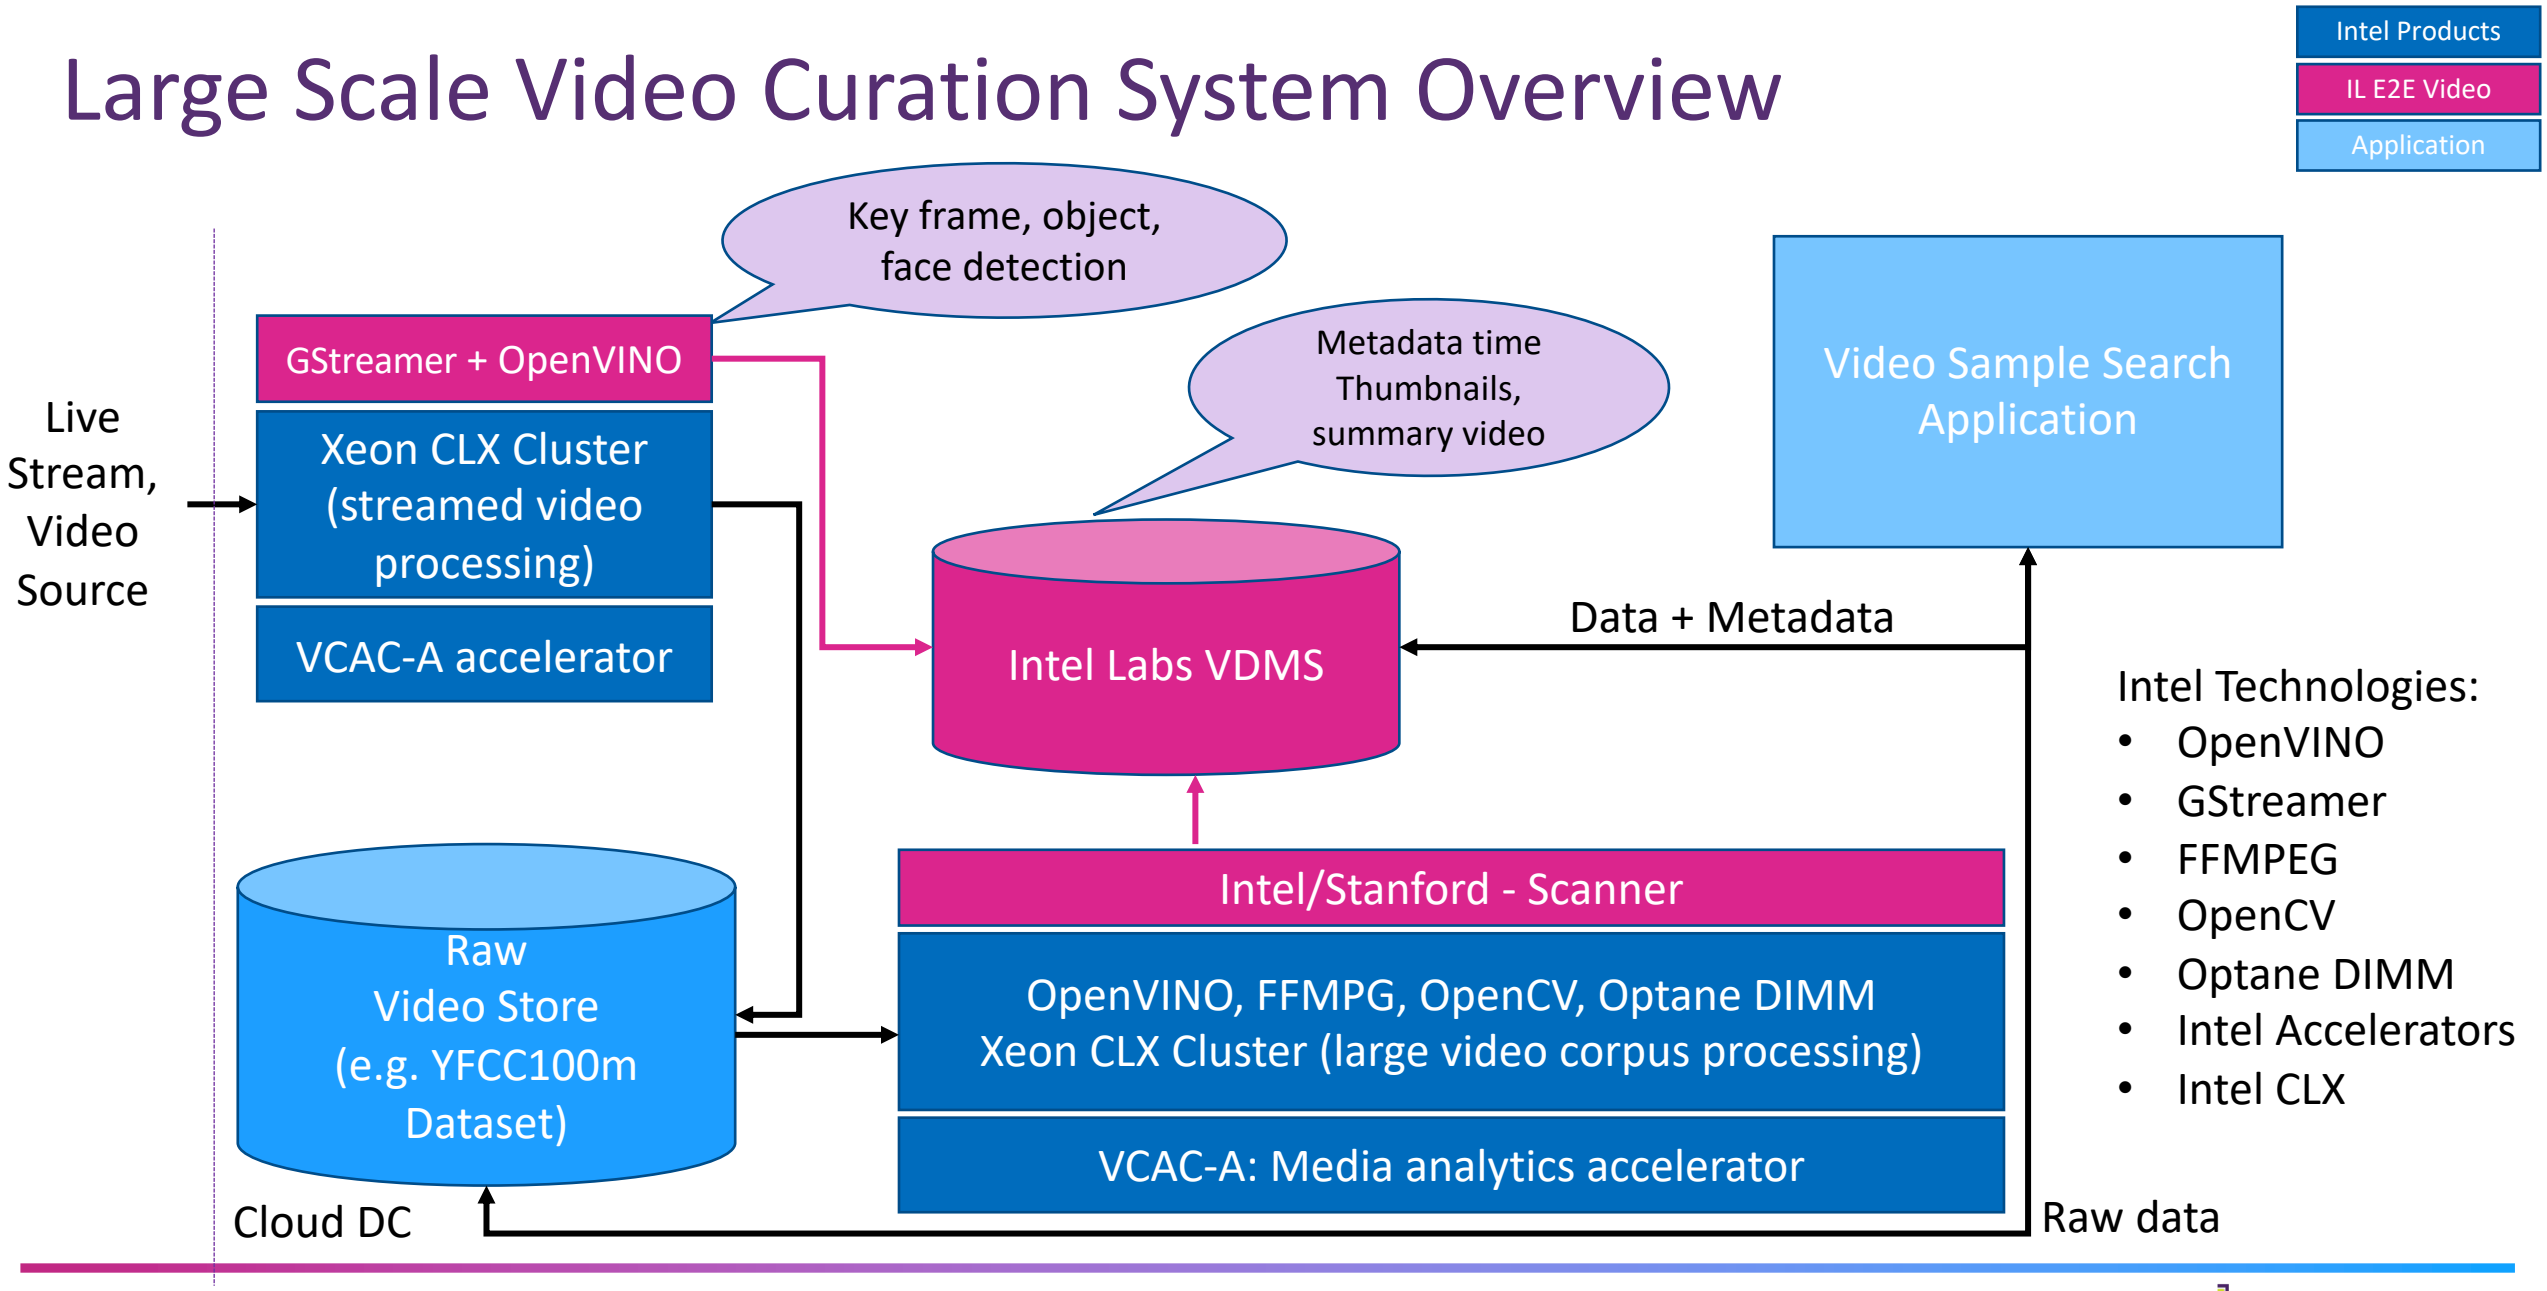


 

# Applied to Applications - Video Curation


Find frames with cars


| | |
|---|---|
| name | 1388383.mp4 |
| time | 0 |
| duration | 30.5 |
| width | 1280 |
| height | 720 |
| fps | 30 |

Summary video where horses with 2 people shows up


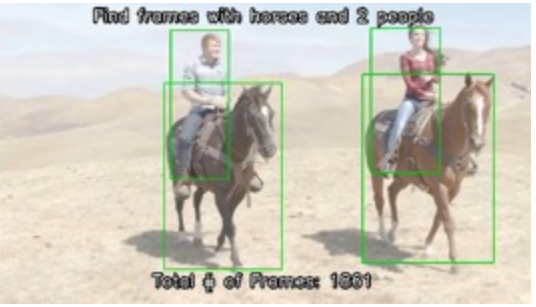


 

# Visual Data Management System

## Visual Workload: Visual Data + Metadata

- Metadata -> Relational Database, Graph Database
- Service for storing the images -> HTTP Server, DFS
- Library for preprocessing -> OpenCV, TensorFlow


Information is segregated among different systems, with different interfaces

## VDMS simplified Interface

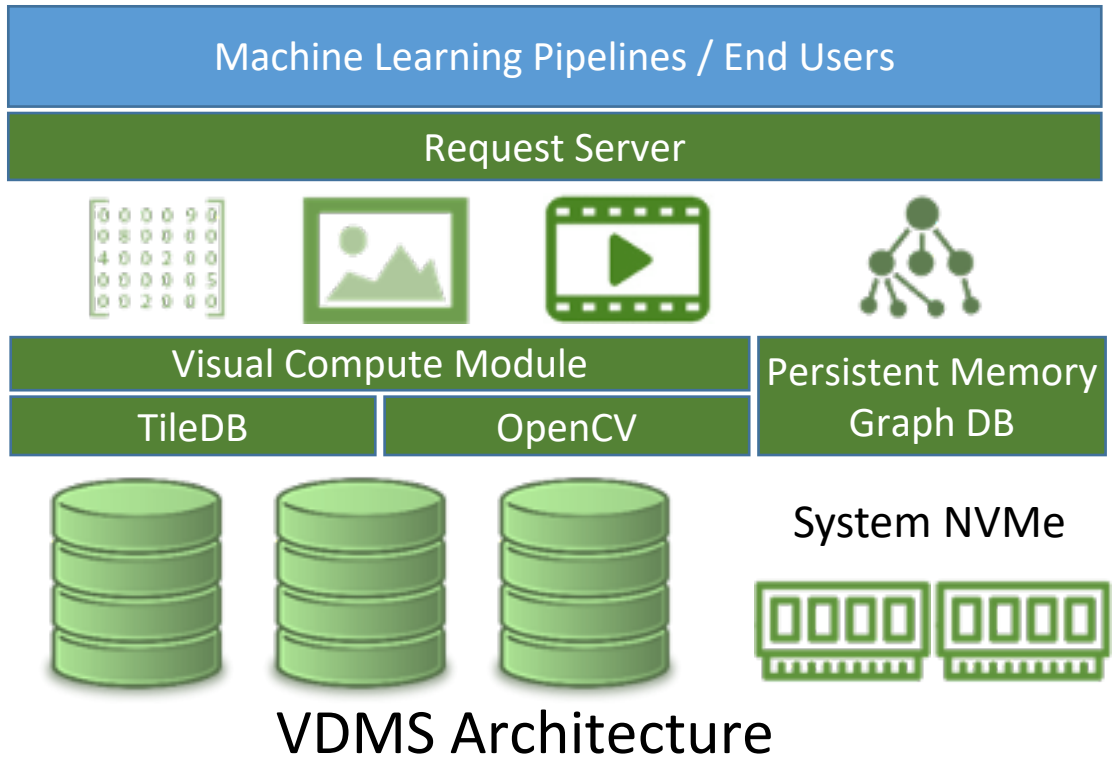


VDMS Architecture

 

# VDMS Query Performance Scaling

(100M images & Over 1B elements – performance evaluation using 5 different queries)

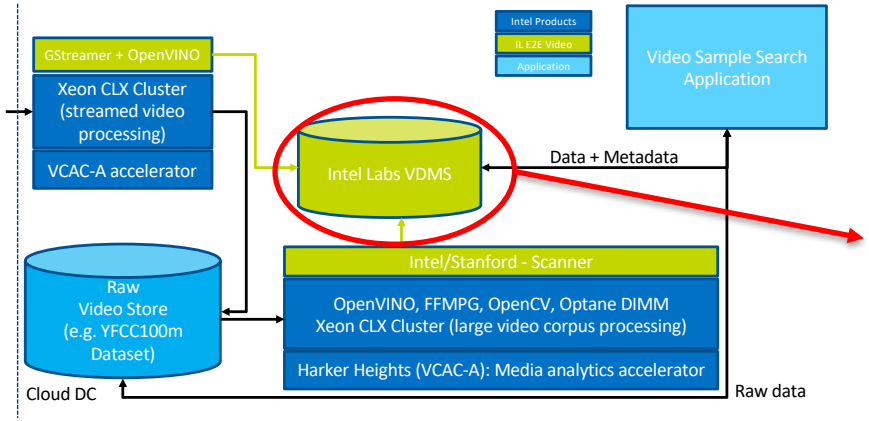


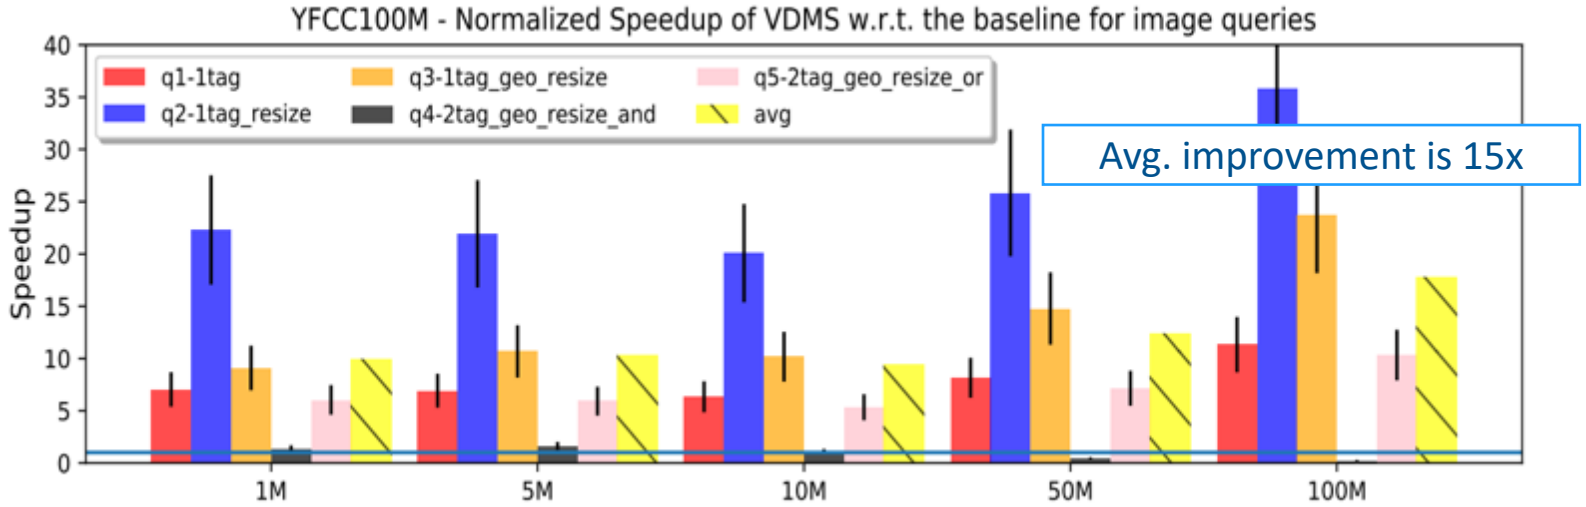


- **VDMS Queries**

1. Find metadata/images with “alligator” tag ***(10x improvement)***
2. Find metadata/images with “alligator” tag and resize to 224x224 ***(35x improvement)***
3. Find metadata/images with “alligator” tag, resize to 224x224, and in a particular geolocation (20 degrees radius) ***(24x improvement)***
4. Find metadata/images with “alligator” AND “lake” tags, resize to 224x224, and in a particular geolocation
5. Find metadata/images with “alligator” OR “lake” tags, resize to 224x224, and in a particular geolocation

**Baseline**: Relational databases (MySQL) + Apache Webserver + OpenCV
**Dataset**: Yahoo Flickr Creative Commons 100 Million Dataset (YFCC100m)
**System**: 2x Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz (Skylake), Ubuntu 16.04

 

# Scanner: Efficient Large-Scale Video Analytics

- **Goal:** Efficient Batch Video Analysis at Scale
- **Other Example Application(s):**
  - Brown Institute for Media Innovation Audiovisual Analysis of 10 Years (200,000 hours) of TV News
  - Facebook Surround 360 Virtual Reality Video stitching across 100's of CPUs.

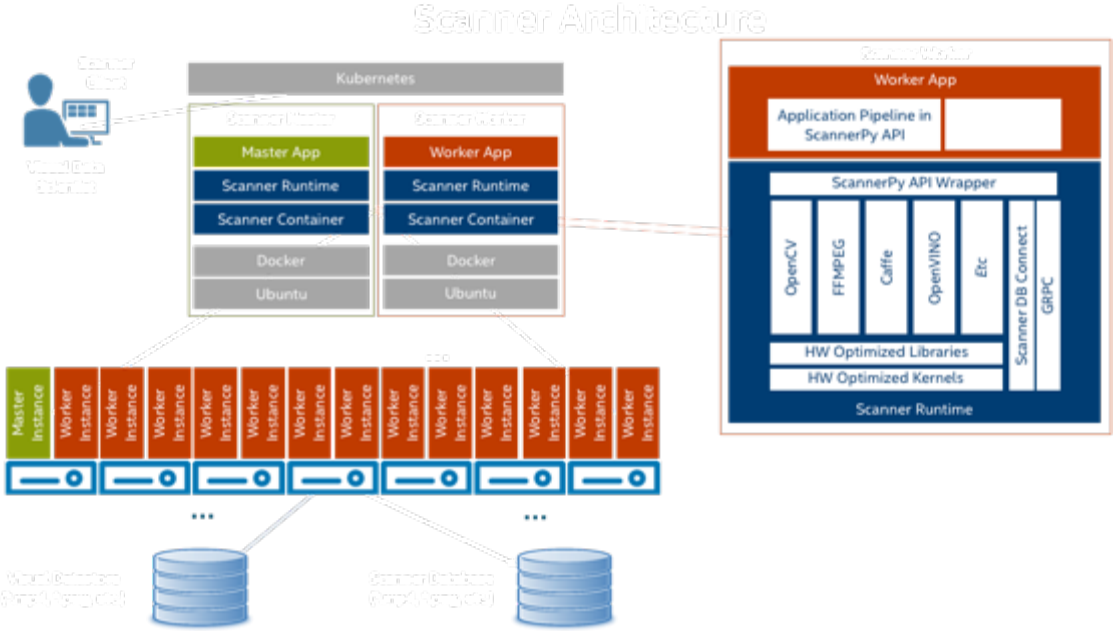


- **Scanner is:**
  - An open source (OSS) video analysis and processing framework. Application developers specify Python pipelines that Scanner distributes across multiple heterogeneous scale-out cloud instances.
- **Current Status:**
  - V3 Open Source Release – April 2019
  - Avail: CPU -- AWS, GCP -- Docker, Kubernetes
  - Optimizations/Libraries: MKL, MKL-DNN, Intel Caffe, Halide, FFMPEG, OpenCV, OpenVINO, Tensorflow, PyTorch, Caffe
  - Presented at Siggraph ’18 and SysML ’19
  - Numerous test applications and pilots conducted

Code: https://github.com/scanner-research/scanner
Docs: http://scanner.run/
Paper: http://graphics.stanford.edu/papers/scanner/scanner_sig18.pdf
Blog: https://itpeernetwork.intel.com/big-video/

Intel ISTC in Visual Cloud Systems

Carnegie Mellon University

 

# Scanner in Action – Stanford Cable TV News Analyzer

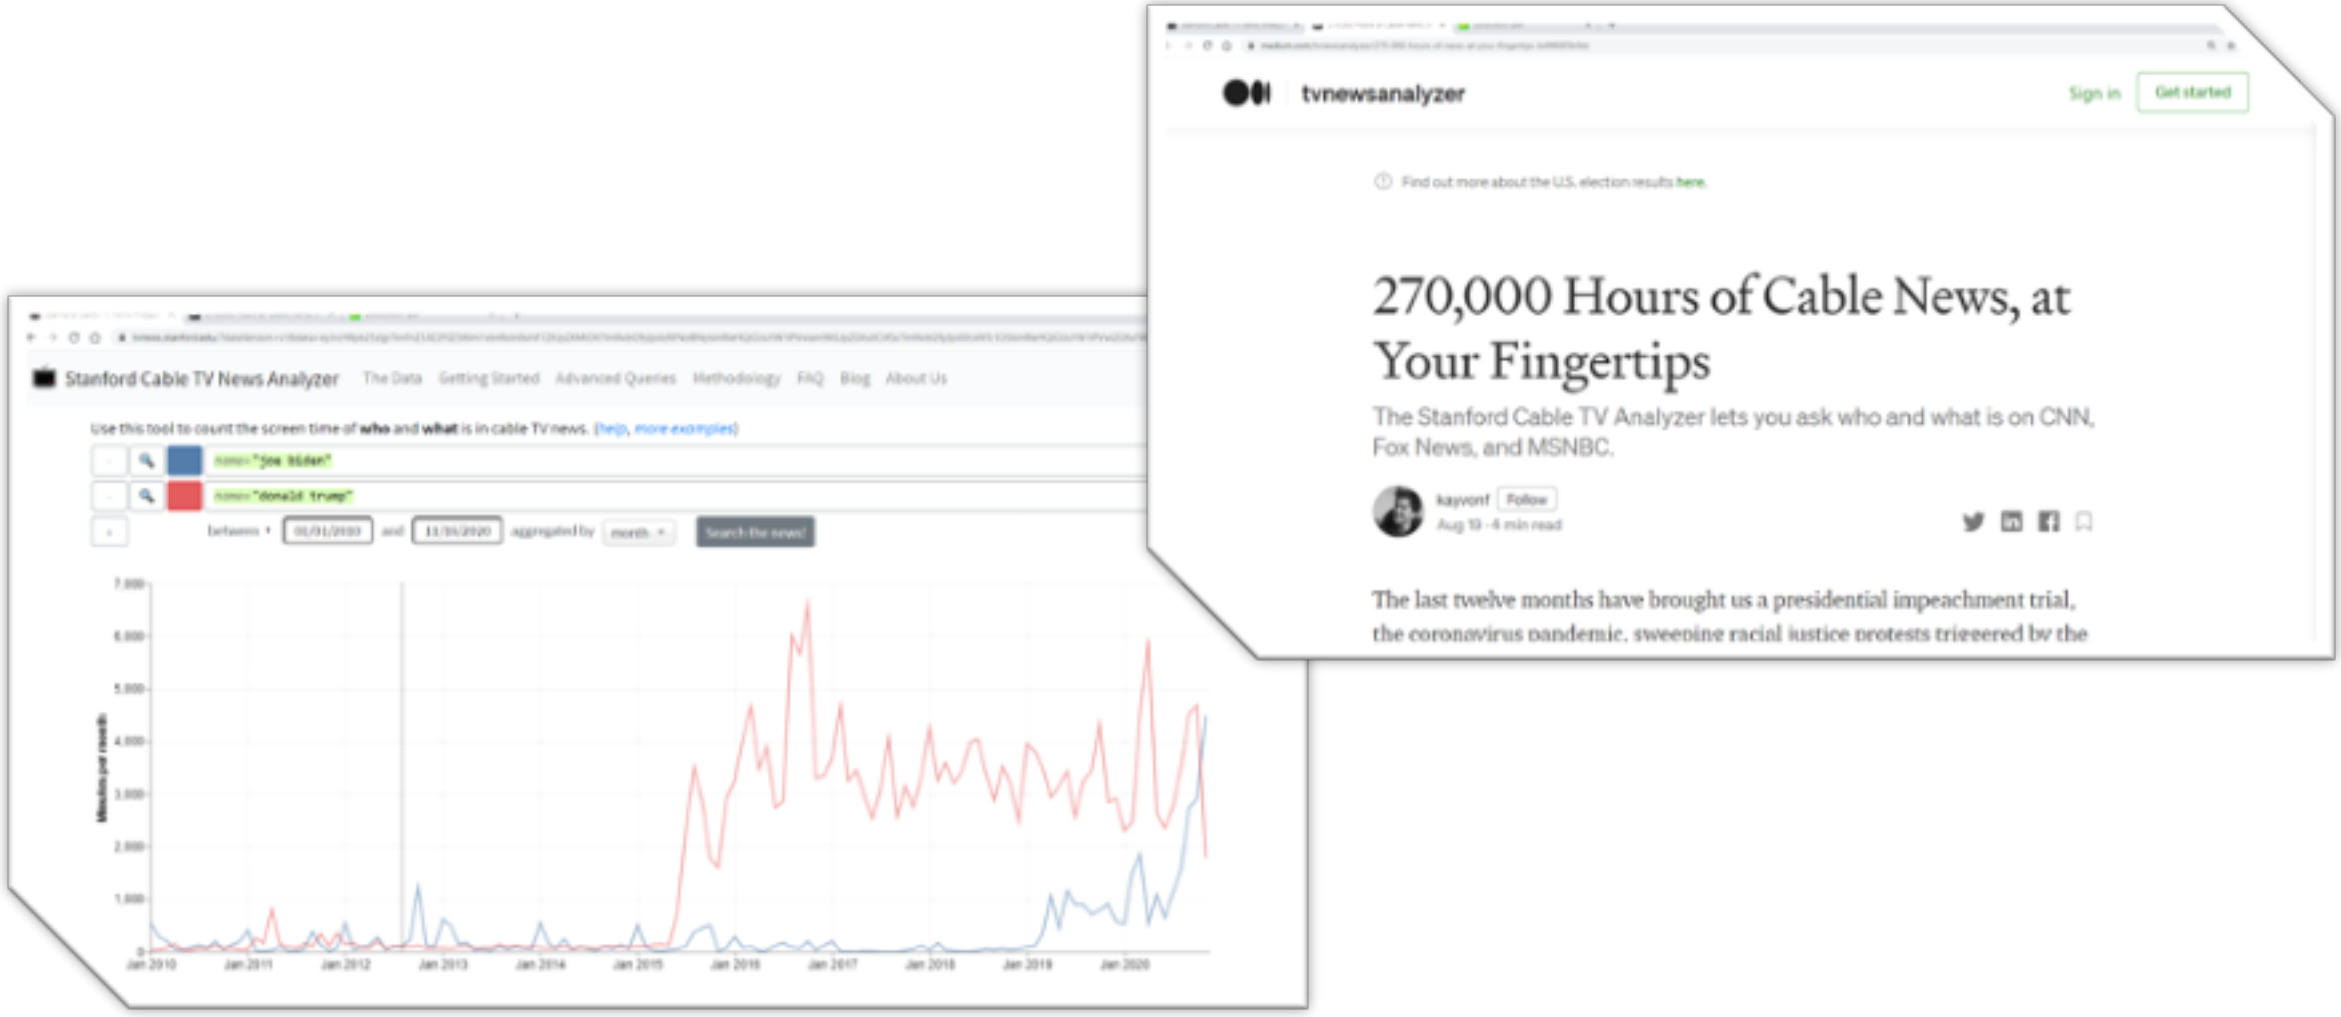


 

# For More Information

| | |
|---|---|
| **Open Visual Cloud** | github.com/OpenVisualCloud 01.org/OpenVisualCloud |
| **Visual Data Management System** | github.com/IntelLabs/vdms/wiki<br>export.arxiv.org/pdf/1810.11832 |
| **Scanner** | github.com/scanner-research/scanner<br>graphics.stanford.edu/papers/scanner/poms18_scanner.pdf |
| **OpenVINO** | openvintotoolkit.org<br>01.org/openvinotoolkit |
| **Stanford Cable TV News Analyzer** | tvnews.stanford.edu |

 

# Video Platform Ecosystem

## It’s more than cameras and tapes now!

 

# Video Platform Ecosystem

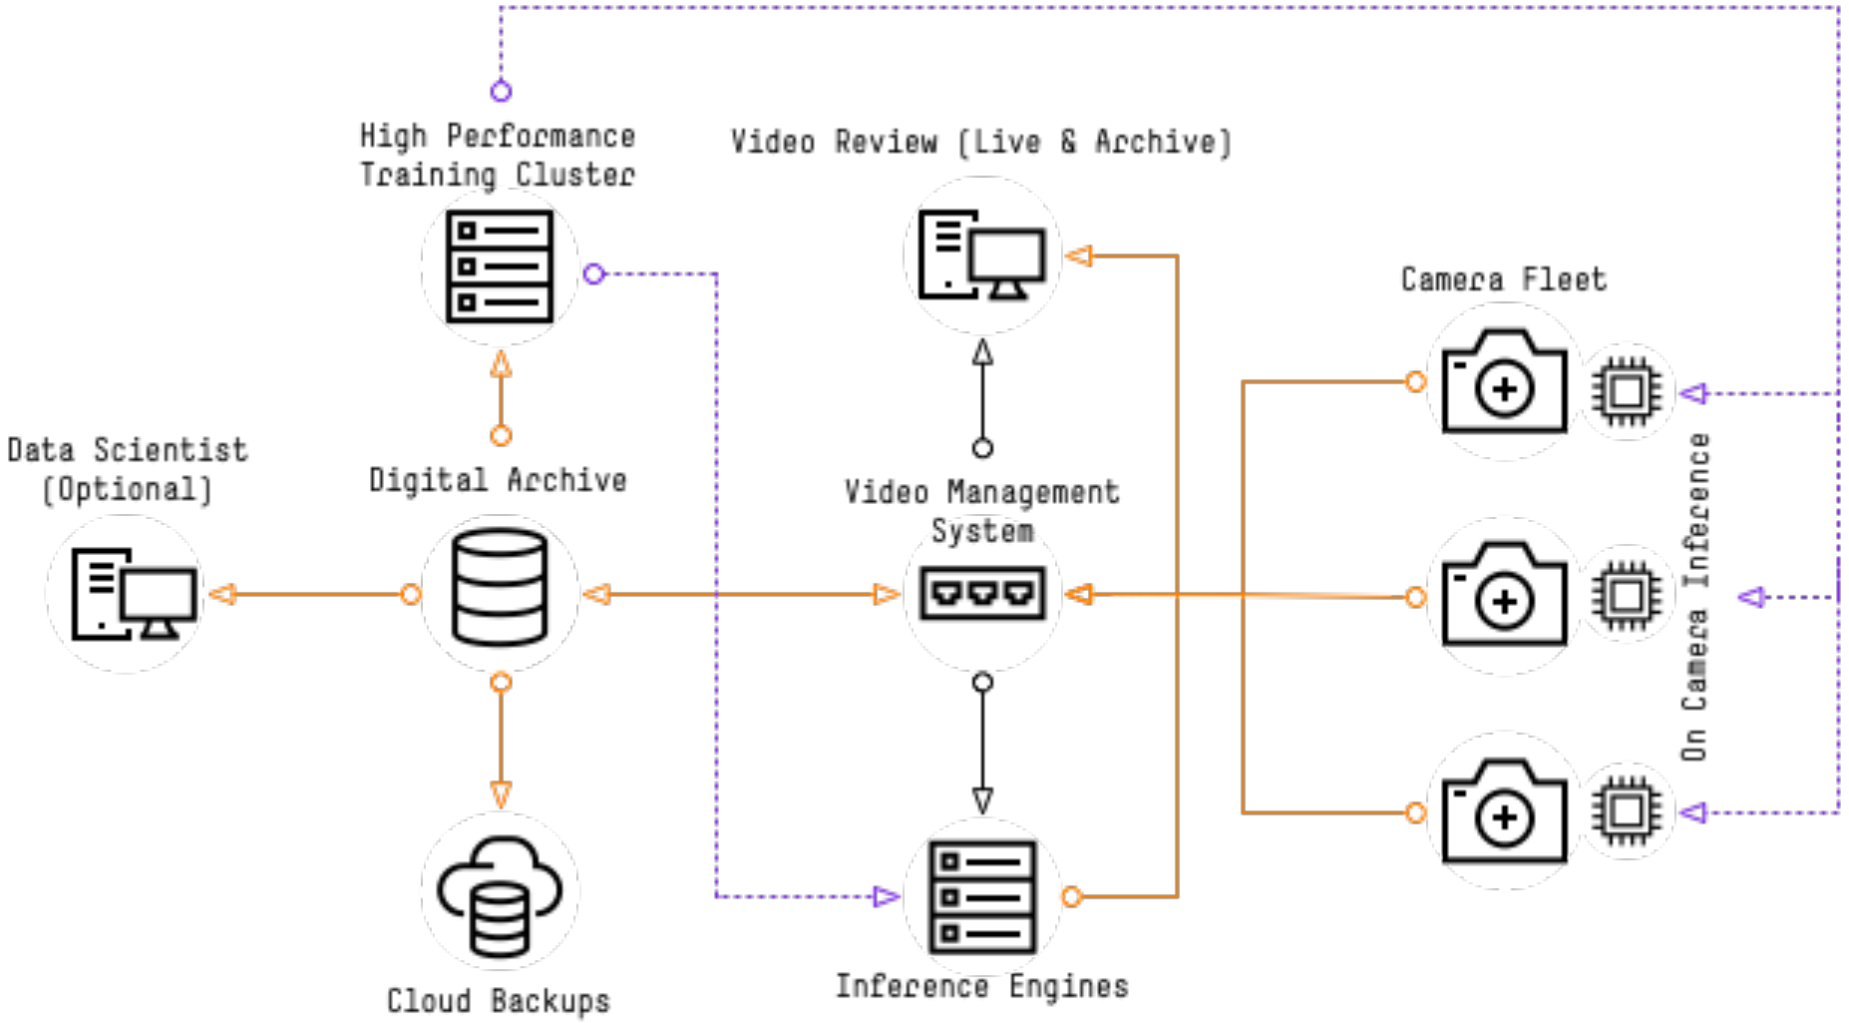


 

# Video Platform Ecosystem

- Inference is now happening at the camera
- Need for more processing power and storage on/near device
- Digital born means better quality, multiple layers but larger data sets
- Multiple uses as we've heard. The archive can feed many other activities, including training different models and analytics
- Governance of both the digital data and the analytics models provides new demands on security and auditing

 

# Security and Governance

## In the analytical age

 

# Security & Governance Implications of Analytics

- Personally Identifiable information
- Facial Detection (using landmarks) vs Facial Recognition
- Trade Secrets may be baked into inference models
- Model and Analytics transparency, why are certain predictions made?
- Capturing data in public but leveraging for new use cases (legislation such as UK Data Protection Act)
- Who can see what? Right to be forgotten (Finding faces can be good!)

 

# Summary

- Video Analytics is providing great value to many organizations
- Leveraging traditional data sets is harder but has value
- Combining video analytics with other analytics and context giving data sources can provide even greater value
- This requires cognitive abilities moving at machine speeds
- The video eco-system is growing and becoming more complex
- Lifecycle management for software and devices is becoming the next big challenge
- Its all pushing data growth even further than before, and looks only to be increasing!

 

# After This Webcast

- Please rate this webcast and provide us with feedback
- This webcast and a copy of the slides will be available at the SNIA Educational Library https://www.snia.org/educational-library
- A Q&A from this webcast will be posted to the SNIA Cloud blog: www.sniacloud.com/
- Follow us on Twitter @SNIACloud

 

# Thank you!

Q & A

 