



Solution White Paper: Alef Stream Processing and Smart Surveillance

Enhance your Stream Processing and Security Solutions with the Software-Defined Mobile Edge

The communications and digital markets are under massive, synchronized transformations. Edge Internet is a distributed Internet architecture combined with a programmable paradigm where computing and connectivity are provided in close proximity to the point of consumption to meet the demands of current and future applications and services.

The synchronized transformations of the communications and digital markets are impacting every trillion-dollar industry in our economy. The next wave of mobile connectivity (5G and next gen wireless networks) and Edge computing is here bringing new applications and services the current infrastructure cannot handle.

“We are at the cusp of a seismic paradigm shift wherein computing and communications will move from a core network with a centralized cloud architecture to the Edge. The reasons are manifold but the basic premise is that in order to serve the data computing and communications demand of objects, sensors, people, resources, compute, and intelligence have to move to the Edge to not only operate in the most cost-effective way but to enable new use cases that can’t be supported by traditional cloud architectures.” – Chetan Sharma, Industry Analyst.

The Edge Internet is a distributed overlay Internet architecture – public and private – with programmable Edge API’s, creating new computing and communications paradigms.

With the rise of 5G and computing environments, AlefEdge is key to unleashing the convergence of networking and computing. AlefEdge’s Software-Defined Mobile Edge platform enables a new set of mobility, cloud and applications APIs that unleash Enterprises and application developers to leverage the full power of the Edge. By utilizing our unique APIs, Enterprises can deploy, orchestrate and manage Edge applications and services.

Table of Contents

Introduction to the Solution	3
Industry Problems	3
Enterprise Problems.....	4
Solution Deep Dive.....	6
How it helps Enterprises.....	8
How Enterprises can Deploy	9

Introduction to the Solution

The Smart Surveillance and Stream Processing Industry is exploding. The ability to process disparate video streams from IoT devices, video surveillance cameras, cameras on machinery etc. to create actionable insights in real-time has created new possibilities for the market. The devices used for the generation of video streams for Smart Surveillance and other Stream Processing use cases are generally regular IP cameras that are continuously generating video streams. Consumer demands are not as stringent as Enterprises demands, therefore, the market needs solutions to meet Enterprise requirements for real-time stream processing.

The Alef Smart Surveillance and Stream Processing solution creates real-time actionable insights based on video uploads while keeping data secure. The solution enables Enterprises to send their video streams to an Edge compute platform with Alef's Stream Processing engine situated at a Micro Edge location to create actionable decisions.


The solution can be leveraged to identify defective products in an assembly line of a manufacturing plant with a camera recording the assembly line and sending video to a Stream Processing platform for real time processing. The video is uploaded to the Edge for processing. Using AI the defective product can be identified in real-time, and the assembly line can be altered or the defective product pulled out of the manufacturing process altogether.

Industry Problems

In recent years, the amount of intelligent CCTV cameras installed in public places for surveillance has increased tremendously, resulting in a large amount of video data produced every moment. This has created a situation where there is an increasing request for the distributed processing of large-scale video data with AI processing to generate real-time actionable insights. In an intelligent video analytics platform, a submitted unstructured video undergoes several multidisciplinary algorithms with the aim of extracting insights and making them searchable and understandable for both human and machine. Video analytics have applications ranging from surveillance to video content management. However, most of the existing solutions rely on a traditional client/server framework to perform face and object recognition while lacking the support for more complex application scenarios. These frameworks are rarely handled in a scalable manner using distributed computing.

The Smart Surveillance and Stream Processing Industry is lacking solutions in three core areas. The first area is a lack of processing and compute power for the devices recording the video streams. The devices are intelligent enough to capture the stream, but there is no local processing or compute on or near the device. This leads to the second issue – no prompt actionable insights are possible with the current architecture. The devices are

required to send their streams to the cloud – creating significant delay and resulting in an inability to create actionable insights. The data being sent to the cloud brings us to the third issue – data storage and compliance. With regulations on data trust and security increasing, it will become more important than ever to have a secure data storage and an audit trail for streams that are processed.



These issues are dominating the Smart Surveillance and Stream Processing industry. The ability to own your data securely, creating actionable insights from localized computing is more important than ever. The Alef Smart Surveillance and Stream Processing solution can solve these industry issues, while focusing on and providing unique services to the Enterprise.

Enterprise Problems

Recently, the volume of video data has increased dramatically on the Internet and various sources are actively contributing toward video generation. YouTube users upload more than 300 hours of videos per minute, almost 58% of downstream traffic on the internet is video, and IntelliVision has deployed more than four million cameras worldwide for surveillance. These unstructured videos are data reservoirs of knowledge and have a direct relation to the real-world events, unlike other data sources. Security agencies process hours of surveillance video collected after some sort of unwanted event occurs to determine moments where suspects might have passed the cameras. Some domains use computer vision techniques for video classification to detect objects, cars, or suspicious behavior in videos. Video data can be turned into useful big data and true video analytics are required to exploit the concepts of computer vision.

With the current abilities of computer vision technology, we have the expertise to mine this visual data to acquire valuable insights about what is happening in the world. Video analytics has stakes from surveillance to space. The demand for intelligent video analytics is expected to be driven by many factors, such as low costs, flexibility, agility and security. As shown in the figure below, the video analytics market is expected to grow from \$4.9B in 2020 to 11.7B in 2025, a 19% CAGR¹. The primary factors contributing to this strong growth in Video Analytics are listed in the figure below:

¹ <https://www.marketsandmarkets.com/Market-Reports/intelligent-video-analytics-market-778.html>

Attractive Opportunities in the Video Analytics Market



Source: Secondary Literature, Expert Interviews, and MarketsandMarkets Analysis

Video processing applications require more time and resources for processing. Extracting insights from any large-scale video analytics system is not easy. Substantial amounts of video data pose a big challenge for video management and manipulation, which require powerful compute clusters to process and mine this video data. Due to large-scale video data, an elastic solution to save and process video data is needed for potential decision making. Most of the existing architectures rely on a traditional client/server framework to perform simple tasks (e.g., face and object recognition) while lacking support for more complex application scenarios (e.g., activity recognition). These frameworks are rarely handled in a scalable manner using distributed computing. Even though intelligent video surveillance has been subject to tremendous advancement, there is still a lack of contributions from the domain of system engineering as well as third party developers.

The cameras installed in Enterprises lack the intelligence required to provide any computing or insight into the stream. They are there only to upload streams to the Cloud. The cameras installed throughout the Enterprise – whether they be for quality control in an assembly line, surveillance or maintenance supervision – have a purpose of generating and transmitting the captured video upstream to the cloud for processing or fault detection or surveillance onsite by employees. The camera has no inherent processing to create any localized insights without any significant compute available nearby. When an Enterprise

may have a given solution *that can do more than video uploads*, the solutions are often niche and not interoperable with other solutions. A manufacturing-oriented quality control solution where the upload video stream is creating real-time actionable insights cannot be easily applied to general surveillance. Enterprises need video analytics solutions which are secure, open and interoperable to solve their video upload and intelligence needs – not one-off solutions.



By implementing the Alef Smart Surveillance and Stream Processing solution, Enterprises can focus on their needs and not worry about the technology that goes into a Video Analytics solution. Alef's solution enables the Enterprise to send one or more video streams to the Stream Processing engine, to be able to apply one or more AI algorithms and send meta data to the cloud for further analytics and insight

Solution Deep Dive

Alef's Stream Processing platform has applications in Industrial IoT and video surveillance. An example is in the Oil/Gas industry. At present the Oil/Gas industry captures video from multiple oil wells every day and uploads this data to a server that is part of a cloud infrastructure. The video stream is then analyzed using an object detection algorithm and its statistics and critical information are generated. This process is time consuming and costly as the oil well has to upload videos streams of massive size, typically in the terabytes, daily.

With Alef's Stream Processing solution running on the Edge on Alef's Software-Defined Mobile Edge (SD-ME) platform, industries do not need to stream their videos onto cloud servers anymore. Video streams will instead be brought to the Edge and processing is completed directly on Alef's local Edge server running its entire software stack. Alef applies object detection algorithms (e.g. Tensorflow AI from Google) on video streams and generates key insights from it. Alef sends all insights in the form of metadata to a central dashboard server. Customers have access to these insights via a dashboard.

Alef's Stream Processing solution also provides an audit feature for the verification of insights generated from the AI engine. Alef stores video files of 1 minute in duration each. These video files can be used for verification of insights generated, should the need arise.

The main components of Alef's Stream Processing solution are:

I. Cloud Dashboard Components

These components are deployed on an EC2 instance in AWS.

a. Upload API

Upload API's is a node server that exposes APIs to store stream insight data and to retrieve them as well.

b. Dashboard API

This component is used for the show insights feature, directly through a web browser.

II. SD-ME Components

All the services listed below are part of the SD-ME framework at an Edge location and are containerized.

a. Node Media Server

The Node Media server accepts the video streams. It is used for publishing live streams. It takes input from an incoming stream and publishes it on live stream.

b. Detection Server

Alef currently uses the Tensorflow Object detection algorithm for real time object detection.

Tensorflow Object detection API is an open source framework built on top of Tensorflow that makes it easy to construct, train and deploy object detection models. It also provides a collection of detection models which are pre-trained on different datasets. The Tensorflow API's object detection component loads frame-by-frame live streaming data to generate insights. These insights/stats are sent to the cloud API server.

c. Audit Server

The audit server is used to ensure accuracy of the generated stats. This module uses FFmpeg to store streaming data in 60-second MP4 files. An audit server starts capturing streaming data as soon as live streaming starts. It will capture videos of 1-minute duration until streaming stops. Audit data will currently be available until the next live streaming begins. After the start of new live streaming, it will override audit data of a previous stream with the latest one.

Alef has kept the Stream Processing solution flexible by segregating all services into different containers keeping in mind if an Enterprise wants to use its own object detection algorithm it can simply run that algorithm in the AI algorithm container.

Alef provides support for only one live stream URL in this first release of the platform. In future releases it will support object detection and data processing on multiple live streams

at the same time. Dashboard support for enabling and disabling audit functionality will be supported in future releases. At present data is captured for audit purposes only. A facility to view data on the dashboard will be added in future releases.

Enterprises can now use their cameras to send video streams of any size to the Edge compute node at the closest Micro Edge Data Center and avail of AI algorithms running on Alef's Stream Processing platform to generate meta data that leads to actionable insights. The video streams can be processed and analyzed at the Edge of the Enterprise network in a cloud native environment at the closest Metro Edge location to the Enterprise. The Alef Smart Surveillance and Stream Processing solution is comprised of APIs which will enable Enterprises to use an AI engine of their choice and to easily obtain insights and analytics via dashboard APIs. The security protocols followed will adhere to Alef's threat model and be an all-encompassing Edge security framework.

Alef Smart Surveillance and Stream Processing can be deployed and leveraged by Enterprises to directly solve their video upload, processing and AI needs, all in a secure SD-ME environment running at the Edge.

How it helps Enterprises

Alef Smart Surveillance and Stream Processing is an easy solution to deploy that will create an open environment for an Enterprise to securely process video uploads and use Alef's Stream Processing platform to obtain real-time actionable insights. By using Alef's open and modular APIs bundled together as Alef Smart Surveillance and Stream Processing, an Enterprise enhances and extends its application capabilities

By connecting cameras already in place within the Enterprise and any new cameras that are implemented to Alef's SD-ME platform, the cameras gain a processing and compute environment in close proximity to the camera. The Edge enables the streams to be processed and intelligence gathered from the AI algorithms running on the SD-ME platform. The insights can range from a security threat within the Enterprise, security authentication to a building or to different locations, ensuring quality control of products and goods being produced, etc. The possibilities are endless when an Enterprise can securely create real-time actionable insights *within its premises* (See Alef AdVision as a reference on secure, real-time actionable insights that can be leveraged to play a targeted ad with low latency from the Edge). Video streams processed at the SD-ME location stay secure within the Edge network thanks to Alef's stringent Edge security protocols and threat-based models implemented.



How Enterprises can Deploy

The Alef Smart Surveillance and Stream Processing solution can be easily deployed for your Enterprise. Signing up, deploying and managing can be done through our Alef EdgeNet Portal. You can learn more about Alef Smart Surveillance and Stream Processing and how to deploy the solution today by learning more in our Resources section.

For more information on any of the documentation, please contact us, use the search bar or contact your sales representative.