Out of the Fog:

Use Case Scenarios

Industry	Smart Cities
Application	Visual Security & Surveillance

Executive Summary

Surveillance and security cameras are being deployed worldwide in record numbers in order to ensure the security and safety of materials, people, and places. These surveillance devices are generating massive amounts of data, with a single camera generating in excess of one terabyte of data per day. Systems of surveillance devices generate data that must be analyzed in real time in order to ensure public safety.

As the number of cameras and the amount of data continues to grow exponentially, new approaches are required to better manage the devices and data. Traditional cloud-based models that were deployed for the lower-resolution cameras in the past, are not ideally suited to these high-definition cameras, due to the massive amount of data, associated latency challenges, network availability and the enormous costs to continually stream the data to the cloud and back.

Fog computing provides the architecture to build cost-effective, real-time and latency-sensitive distributed surveillance systems that help to preserve privacy challenges in open environments.

Challenges

- The cloud can’t scale to support wide-scale surveillance in applications such as buildings, highways, cities, airports, etc.
- Rapid security decisions must be made on location in real time; latency delays poise public safety risks.
- High definition cameras generate terabytes of data per day, which is unmanageable by the cloud alone.
- Video surveillance requires highly-reliable connectivity and bandwidth, which can be challenged in emergency situations.

Solution

- Fog-based deployments provide the opportunity to build real-time, latency-sensitive distributed surveillance and analytics.
- Fog computing enables real-time tracking, anomaly detection and insights from data collected over long time intervals.
- Heterogeneous processor capabilities of fog runs parts of the video analytics algorithm on conventional processors or accelerators.

Technology

- Fog nodes intelligently partition video processing between cameras and the cloud.
- Fog provides the secure transportation and storage of data.
- Video analytics algorithms can be located on fog nodes close to the cameras, for rapid latency and reaction to public safety situations.

Visual Security and Surveillance

Surveillance and security cameras are being deployed worldwide. These cameras are used to ensure security of materials, people, and places. In addition, these cameras have the ability to generate a massive amount of data, which can exceed terabytes per day for a single camera.

Traditional cloud models that were deployed for low-resolution cameras aren’t scalable with the new 1080p and 4K cameras because of the sheer availability and/or cost of network transport. While much of the data generated is mundane, abnormal events need to be detected rapidly. Decisions on security need to be made on location, where the data is capture, and cannot be made solely in the cloud.

Machine vision is also a prime candidate for accelerators and dynamic updating of various algorithms in both hardware and software. These cameras are capturing images of people, places, or things and are tightly coupled to decision–making, which requires a heightened level of security of the camera’s software and hardware assets.

Smart cities, smart homes, building complexes, retail stores, public transportation, manufacturing and enterprises increasingly rely on surveillance cameras to monitor assets, identify unauthorized access, and increase the safety of its inhabitants and property. The sheer bandwidth of visual (and other sensor) data being collected over a large-scale network makes it impractical to transport all the data back and forth to the cloud to obtain real-time insights.

A particularly demanding application is surveillance of areas with many people and objects moving through them, such as in a crowded city or at an airport.

City-scale deployments that include placing cameras on traffic lights and in remote areas don't have high-bandwidth connectivity to the cloud to upload the collected video. Even if the video could fit over the network infrastructure, consideration must be gigven to the cost of the transmission and the need for high network availability. Real-time monitoring and detection of anomalies (intruders into a building, the fall of an elderly citizen, the misfiring of a piece of manufacturing equipment) pose strict low latency requirements on surveillance systems. Timeliness is important from the standpoint of both detection and response.

Additionally, privacy and security concerns must be addressed when using a camera as a sensor that collects image data so that the images do not reveal a person's identity or reveal confidential contextual information to any unauthorized parties. For example, a manufacturing plant may not want to transmit images containing intellectual property to the cloud for security reasons.

Key Features Enabled by Fog Computing

Fog computing is the horizontal architecture designed to build real-time, latency-sensitive distributed surveillance systems that maintain privacy.

Fog computing provides a mesh of fog nodes to intelligently partition video processing between other fog nodes co-located with cameras and the cloud so as to enable real-time tracking, anomaly detection, and insights from data collected over long time intervals.

Video analytics algorithms can be located on fog nodes close to the cameras, and take advantage of the heterogeneous processor capability of fog, running parts of the video analytics algorithm on conventional processors or accelerators.

An End-to-End Use Case: Airport Visual Security

Airport visual security, called surveillance, illustrates the complex, data-intensive demands required for real-time information collection, sharing, analysis, and action.

It starts with a look at the passenger’s journey:

- Leaves home and drives to the airport
- Parks in the long-term parking garage at the airport
- Takes bags to airport security checkpoint
- Checks in self and bags, which are scanned and checked in
- Goes through security and proceeds to boarding gate
- Retrieves bags from conveyer belt at destination airport
- Proceeds to rental car agency; leaves airport

In the vast majority of cases, this travel scenario is without incident. But when one or more threats are present, the visual security requirements become infinitely more complicated. For example:

- The vehicle entering the airport is stolen
- The passenger's name is on a no-fly list
- The passenger leaves his luggage unattended someplace in the airport
- The passenger's luggage doesn't arrive with the flight
- The luggage is scanned and loaded on the plane, but it is not picked up by the correct passenger
- An imposter steals or switches a boarding pass with another passenger and gets on someone else's flight
- The passenger takes someone else's luggage at the arrival terminal

Catching these possible threats requires an extensive network of several hundred surveillance cameras across the outbound and inbound airports. This involves massive amounts of data: Approximately one terabyte of data per camera per day must be transmitted to security personnel or forwarded to local machines for scanning and analysis.

In addition, the information must be shared with law enforcement agencies, who will need data originating from multiple systems about the suspect passenger's trip, from the point of origination to arrival. Finally, all of the video and data must be integrated with a real-time threat assessment and remediation system.

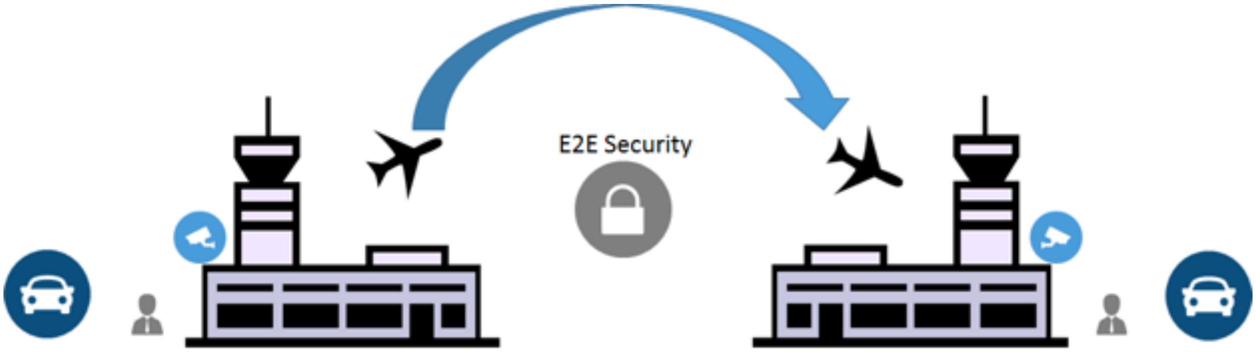

Cloud and Edge Approaches. In an edge-to-cloud design, every camera (edge device) in the airport transmits directly to the cloud for processing, as well as the other relevant data collected from the passenger's travel records.

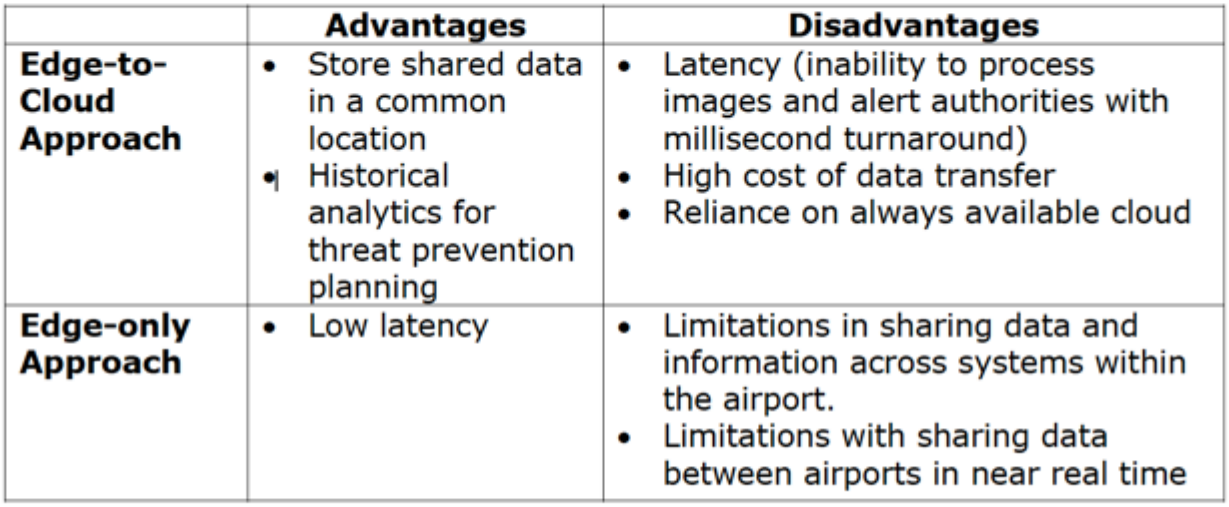

	Advantages	Disadvantages
Edge-to-Cloud Approach	• Store shared data in a common location • Historical analytics for threat prevention planning	• Latency (inability to process images and alert authorities with millisecond turnaround) • High cost of data transfer • Reliance on always available cloud
Edge-only Approach	• Low latency	• Limitations in sharing data and information across systems within the airport. • Limitations with sharing data between airports in near real time

While there are advantages to both approaches, the disadvantages can lead the systems susceptible to incidents.

Fog Computing Approach. The fog computing approach greatly reduces the latency, network bandwidth and availability constraints. The architecture for fog computing is based on eight pillars of elements, as identified by the OpenFog Consortium. In an airport, the following pillars are addressed:

- **Security**: The airport visual security scenario is a physically distributed fog deployment. Thus, physical possession is in scope for the security analysis. Transportation and storage of data must also be secure as much of the data which may contain personally identifiable information.

- **Scalability**: The fog system must adapt with the business needs as it relates to system cost and performance. When a new airport terminal, gate, or additional sensors and equipment are added, the solution must scale and not require a completely new deployment.

- **Open**: Openness is essential for the success of a ubiquitous fog computing ecosystem for IoT or 5G platforms and applications. Proprietary or single vendor solutions can result in limited supplier diversity, which can have a negative impact on system cost, quality and innovation.

- **Reliability/Availability/Serviceability**: The entire surveillance system must be reliable, available, and serviceable which includes orchestration of existing or new resources. As new object recognition models are trained for visual analytics, these inference engine models should be updated on or near edge devices without down time to the system.

- **Programmability**: Programming at the hardware level may be necessary for the visual analysis in order to perform inference on images.

Architectural View of Fog in Airport Surveillance

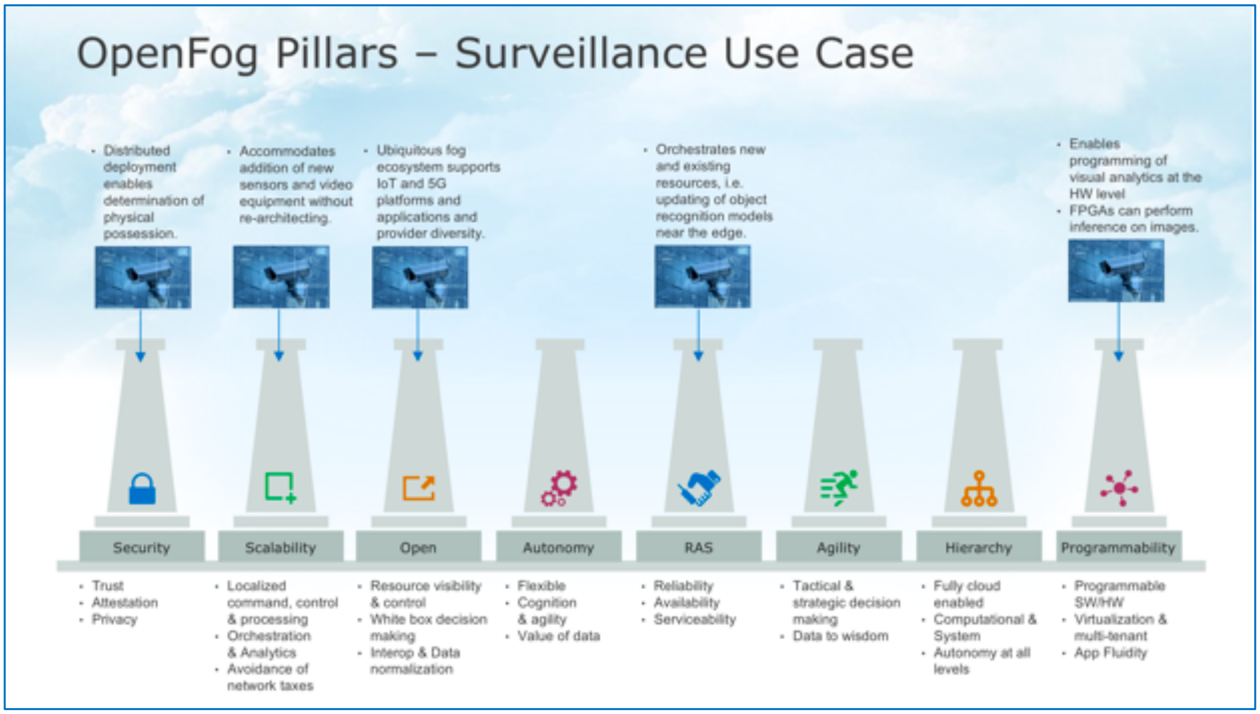

OpenFog Pillars – Surveillance Use Case

- Distributed deployment enables determination of physical possession.
- Accommodates addition of new sensors and video equipment without re-architecting.
- Ubiquitous fog ecosystem supports IoT and 5G platforms and applications and provider diversity.
- Orchestrates new and existing resources, i.e. updating of object recognition models near the edge.
- Enables programming of visual analytics at the HW level
- FPGAs can perform inference on images.

Security
- Trust
- Attestation
- Privacy

Scalability
- Localized command, control & processing
- Orchestration & Analytics
- Avoidance of network taxes

Open
- Resource visibility & control
- White box decision making
- Interop & Data normalization

Autonomy
- Flexible
- Cognition & agility
- Value of data

RAS
- Reliability
- Availability
- Serviceability

Agility
- Tactical & strategic decision making
- Data to wisdom

Hierarchy
- Fully cloud enabled
- Computational & System
- Autonomy at all levels

Programmability
- Programmable SW/HW
- Virtualization & multi-tenant
- App Fluidity

What is Fog Computing?

Fog computing is a system-level horizontal architecture that distributes resources and services of computing, storage, control and networking anywhere along the continuum from Cloud to Things.

- **Horizontal architecture:** Supports multiple industry verticals and application domains, delivering intelligence and services to users and business.
- **Cloud-to-Thing continuum of services:** Enables services and applications to be distributed closer to things, and anywhere along the continuum between Cloud and Things.

- **System-level:** Extends from the Things, over the network edges, through the Cloud, and across multiple protocol layers – not just radio systems, not just a specific protocol layer—not just at one part of an end-to-end system, but a system spanning between the Things and the Cloud.

About the OpenFog Consortium

Video surveillance is just one of many industry use cases whose commercial viability will depend on fog computing in order to achieve the rapid response, bandwidth and communication necessary in advanced digital applications.

The OpenFog Consortium is a global nonprofit formed to accelerate the adoption of fog computing in order to solve the bandwidth, latency, communications and security challenges associated with IoT, 5G and artificial intelligence. Our work is centered around creating a framework for efficient and reliable networks and intelligent endpoints combined with identifiable, secure, and privacy-friendly information flows in the Cloud-to-Things continuum based on open standard technologies. For more information, please contact us at info@OpenFogConsortium.org.