WHITE PAPER

Security
High Value Asset Surveillance

Wipro Provides High Value Asset Surveillance Edge Analytics

Wipro High Value Asset Surveillance (HVAS) solution provides edge analytics to enable real-time response to crime or theft of high value assets (HVA). The solution utilizes OpenNESS.

Introduction

Enterprises in retail, banking, health care, airports, and other industries deploy a variety of security solutions to protect high value assets (HVA). HVAs include expensive jewels, automated teller machines (ATMs), cash on premises, vaults and safety deposit boxes, expensive medical equipment, or even police stations.

HVA protection solutions include security personnel and multiple layers or levels of security and surveillance technology, including sensor detection or video or drone monitoring. Some of the challenges that come with this increased security include the cost of backhauling data traffic, the ongoing operating expenses of these security technologies, and human lives being placed at risk.

Traditionally, the HVA geofencing application is run in a cloud network, where a cloud server performs high-precision video analytics using high-quality video traffic that is sent to this server via the network. The roundtrip delay to this central server function and back, though, adds to latency. This processing must happen very quickly for a fast response to a potential incident. In many cases, the backhaul traffic over internet links lack service level agreements (SLAs), which can impact network predictability.

One new addition to this solution set that overcomes some of these challenges is HVA surveillance (HVAS), which adds the ability to process and analyze surveillance video for risks or threats and send alert notifications in real time to a surveillance center. HVAS solutions monitor the video feed of a surveillance camera via a video analytics engine (VAE). The analytics engine will identify the object, compare it with high value assets classified in the database, and then apply appropriate geofencing parameters that define the criteria for raising an alert. When the object is moved out of this geofenced area, an alert is raised.

The answer is to move the HVAS analytics out of the cloud and nearer to the HVA. To do this, Wipro, an Intel® Network Builders ecosystem member and Intel® Network Builders Edge Partner, has developed an HVAS solution that utilizes multi-access edge computing (MEC) servers with Open Network Edge Services Software (OpenNESS) to deploy and manage the onboarding of the video analysis application on edge servers. OpenNESS delivers application lifecycle management on different edge servers, either on-premises or at the network edge.

The Wipro HVAS utilizes an SD-WAN end-to-end quality of service (QoS) feature provided by Lavelle Networks, ScaleAOn, which delivers critical notifications generated by the HVAS video analytics application to the security monitoring center.

The service orchestration capability of OpenNESS (see Figure 1) enables deployment of the solution at scale by orchestrating with multiple edge instances of the service based on the need. A learning model obtained from one node can

Table of Contents

be propagated across all instances of the service. The solution also provides assured quality of the video feed by leveraging the traffic steering capability of SD-WAN.

OpenNESS Overview

Wipro uses OpenNESS to add orchestration features to its network edge–deployed software. OpenNESS is an open source software toolkit that enables multi-access edge computing. OpenNESS is an open source software toolkit that enables highly optimized and performant edge platforms to onboard and manage applications and services with cloud-like agility across any type of network.

Some of the key challenges that are mitigated by OpenNESS in the Wipro HVAS application include the following:

- **Support for multiple access technologies:** Works with 5G, LTE, Wi-Fi, and wired networks.
- **Edge orchestration:** Exposes northbound APIs that a central orchestrator such as ONAP can use to federate edge orchestration.
- **Deployment:** Can be implemented at either the on-premises edge or the network edge.
- **Hardware abstraction:** Supports a template for resource description that simplifies deployment.

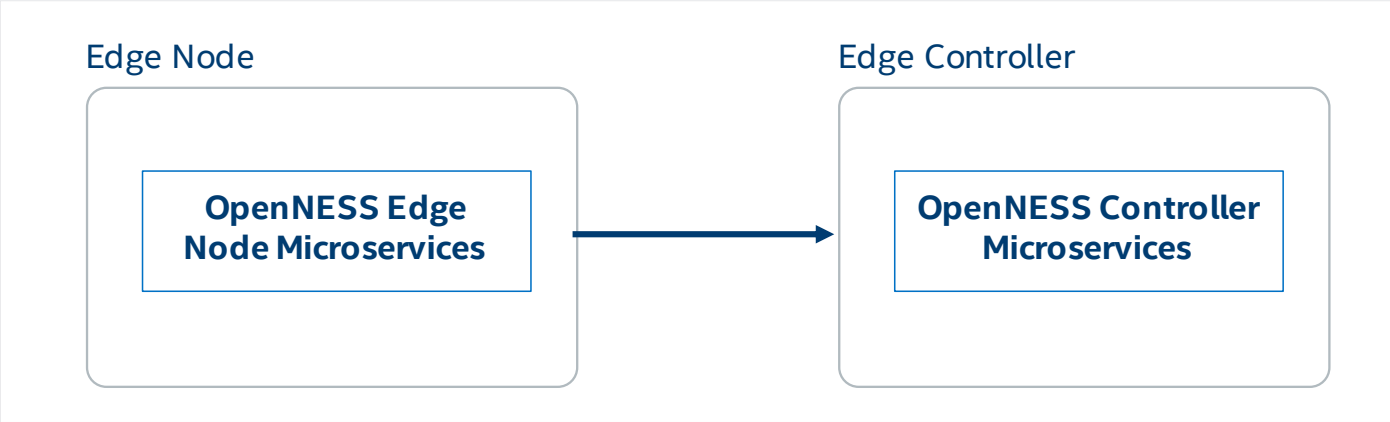

Figure 1. OpenNESS overview.

An OpenNESS subsystem consists of one or more OpenNESS edge nodes and a controller node. Both nodes host specific microservices, arranging an application's requirements as a collection of independently deployed services.

An edge node hosts a set of OpenNESS microservices, edge compute applications, and network functions. OpenNESS edge node microservices enable execution of edge compute applications natively on the edge node, or forward the user traffic to applications running on platforms connected to the edge node on a local breakout. The OpenNESS edge node runs on a real-time kernel and leverages the open source Data Plane Development Kit (DPDK) to accelerate the data plane implementation.

The OpenNESS controller consists of a set of microservices that enable existing edge compute cloud orchestration and application lifecycle management. The OpenNESS controller may be hosted locally, or be hosted in an enterprise or public cloud to manage edge nodes.

An OpenNESS application can be categorized as follows depending on the servicing of end user traffic:

- Producer application: Provides services to other applications running on the edge compute platform. Producer applications do not serve end user traffic directly.
- Consumer application: Serves end users traffic directly. Consumer applications may or may not subscribe to the services from other producer applications on the edge node.

Wipro's HVAS solution services the end user traffic and hence uses the consumer application deployment option of OpenNESS. This producer/consumer categorization allows a service subscription-based application architecture.

OpenNESS supports two environments for building an edge platform. The first one is based on a Kubernetes environment, where the controller services are run as part of the Kubernetes master, while the edge node's services run in the cluster worker node. The second is based on a KVM and native Docker run-time environment for the edge node, and the controller runs in a separate node. The HVAS service use case uses the latter environment to deploy and manage the edge applications.

Overview of High Value Asset Surveillance by Wipro

The Wipro HVAS video analytics solution consists of three functions:

- Training data set generation
- Data preprocessing
- Training of the convolutional neural network (CNN) model

In the **data set generation phase**, a machine learning (ML) model is used to perform object detection. The data on which the ML model is trained is vital, playing an important role in the accuracy of the model. Wipro uses the Open Source Computer Vision Library (OpenCV) to extract frames from a video stream to train the system to understand particular threats to the HVA.

In the **data preprocessing function**, Wipro uses the data preprocessing and data augmentation module of the open source Keras neural-network library. Using the Keras library, training data that has preprocessed with grayscale is obtained, the image data is compressed, and the dataset is prepared for "training model" processing.

The last component of the process is **training the CNN model**. The Wipro CNN model architecture is designed for a binary classification of objects detected. Wipro has trained the CNN model to determine the presence of an HVA.

The dataset is then split into a training data set and test data set for validating the performance of the CNN model. The F1 score is the accuracy classification metric considered for analyzing the performance of the model. The hyper-parameters are chosen accordingly to get accurate predictions from the CNN model, and the model is saved for later usage.

As seen in Figure 2, a surveillance camera focused at the HVA sends its video stream to the HVAS video analytics application. This application is hosted on a Docker container and is onboarded to an OpenNESS edge node. The HVAS video analytics application is trained to monitor the presence of HVAs under surveillance. Once the object is removed from the vicinity of the camera, an alert message is triggered by the analytics engine and is sent to the SD-WAN controller for dynamic policy enforcement. The SD-WAN controller pushes traffic policy to the backhaul network, giving critical priority to the alert message.

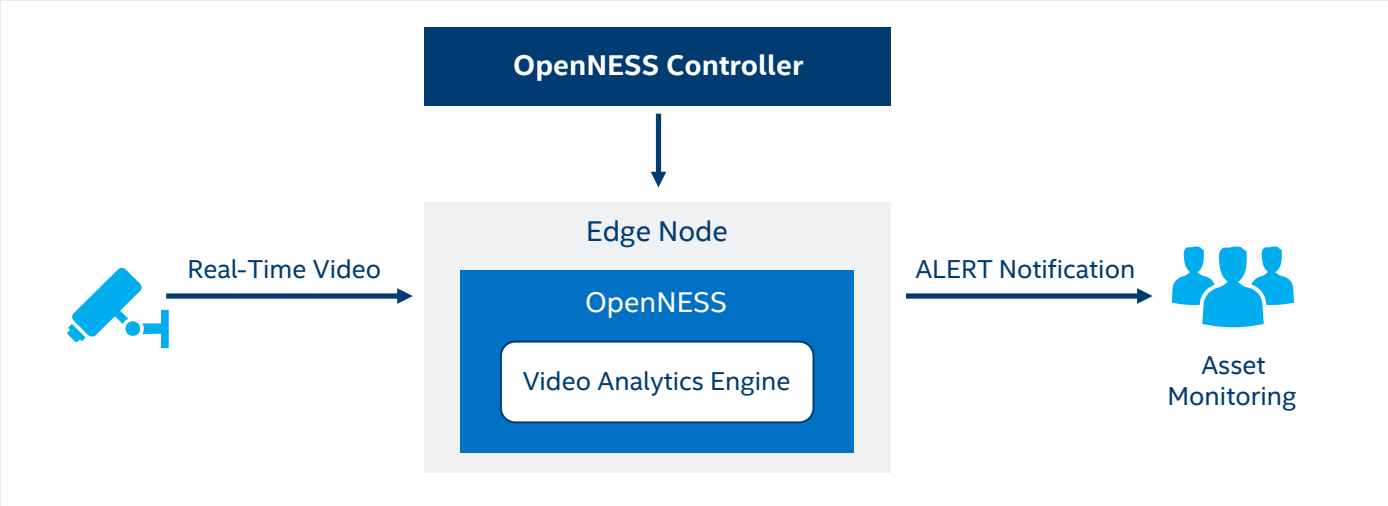

Figure 2. Wipro's Solution for High Value Asset Surveillance.

SD-WAN

Wipro's solution is enhanced by the addition of Lavelle Networks' ScaleAOn SD-WAN platform. ScaleAOn is flexible because it runs on a wide range of hardware form factors and supports a wide variety of next-generation edge networking scenarios.

The SD-WAN edge software features zero-touch network connectivity to mobile backhaul networks. Traffic from the edge surveillance cameras is carried over dynamic encrypted tunnels. These tunnels are set up on demand and are configured with the highest traffic prioritization to ensure that intelligent edge notifications are sent even during heavy bandwidth utilization in the 4G or 5G coverage area.

The policies in the ScaleAOn SD-WAN solution are intent based, where a named network object (like LAN or WAN) or a network direction (like LAN to WAN) is used instead of network parameters like IP address or interface name. This intent-based policy makes it easy to connect a host of devices, like HVAS cameras, and transport their processed information to the cloud data center or on-premises surveillance servers.

Intent-based policies will treat notification traffic as critical with the highest priority. When an alert notification message is received, the QoS policy will be pushed to all devices in the path between the edge node and the destination. Traffic will be immediately forwarded, and normal traffic will be buffered until the critical traffic is pushed through completely.

End-to-End Solution Architecture

The integration of OpenNESS into the HVAS solution provides edge video analytics deployment as shown in Figure 3.

The HVAS application is packaged as a VM. This VM is then onboarded into the OpenNESS edge node via the OpenNESS controller. The upstream interface is connected to the video feed. The downstream interface is connected to the SD-WAN device. The SD-WAN controller deployed in the cloud will push QoS policies for traffic prioritization. The web UI at the remote spoke of the SD-WAN device will display these notifications.

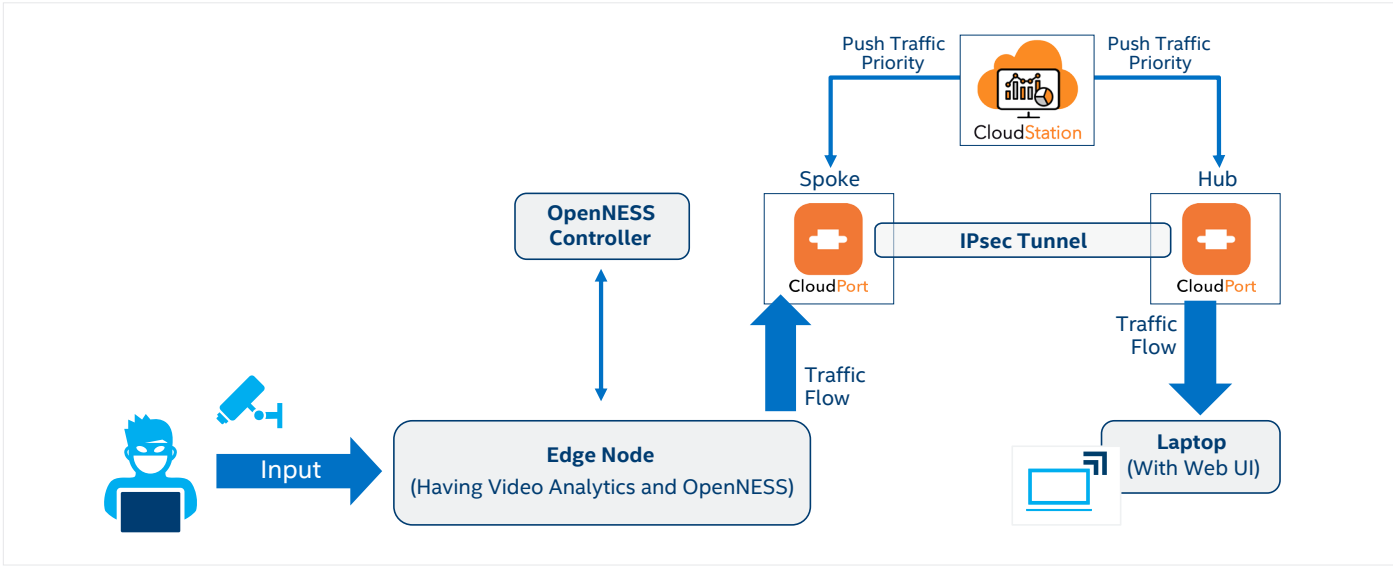

Figure 3. Block diagram of SD-WAN Integration with HVAS solution.

As shown in Figure 4, the edge node has four functions used for HVAS implementation:

1. **OpenNESS Data Plane Services:** This steers traffic toward applications running on the edge node or the local breakout port. Traffic policies are configured on the OpenNESS edge controller and pushed to the data plane services, such that traffic steering is applied to either redirect the traffic to edge applications for further analysis or pass the packets through the downstream interface to the packet core for traffic forwarding over SD-WAN.
2. **OpenNESS Enhanced Platform Awareness Microservices:** These microservices include edge authentication agent (EAA), edge virtualization agent (EVA), edge lifecycle agent (ELA), syslog, DNS, and others. These microservices manage application lifecycle, DNS resolution, application enrollment, and more.
3. **Evolved Packet Core (EPC):** The edge node is attached to the SGi interface of an EPC. Traffic from the EPC arrives as IP traffic and is steered as needed to edge applications. The EPC combines both user and control plane.
4. **HVAS Application:** Video analytics engine running in a VM or as a Docker container.

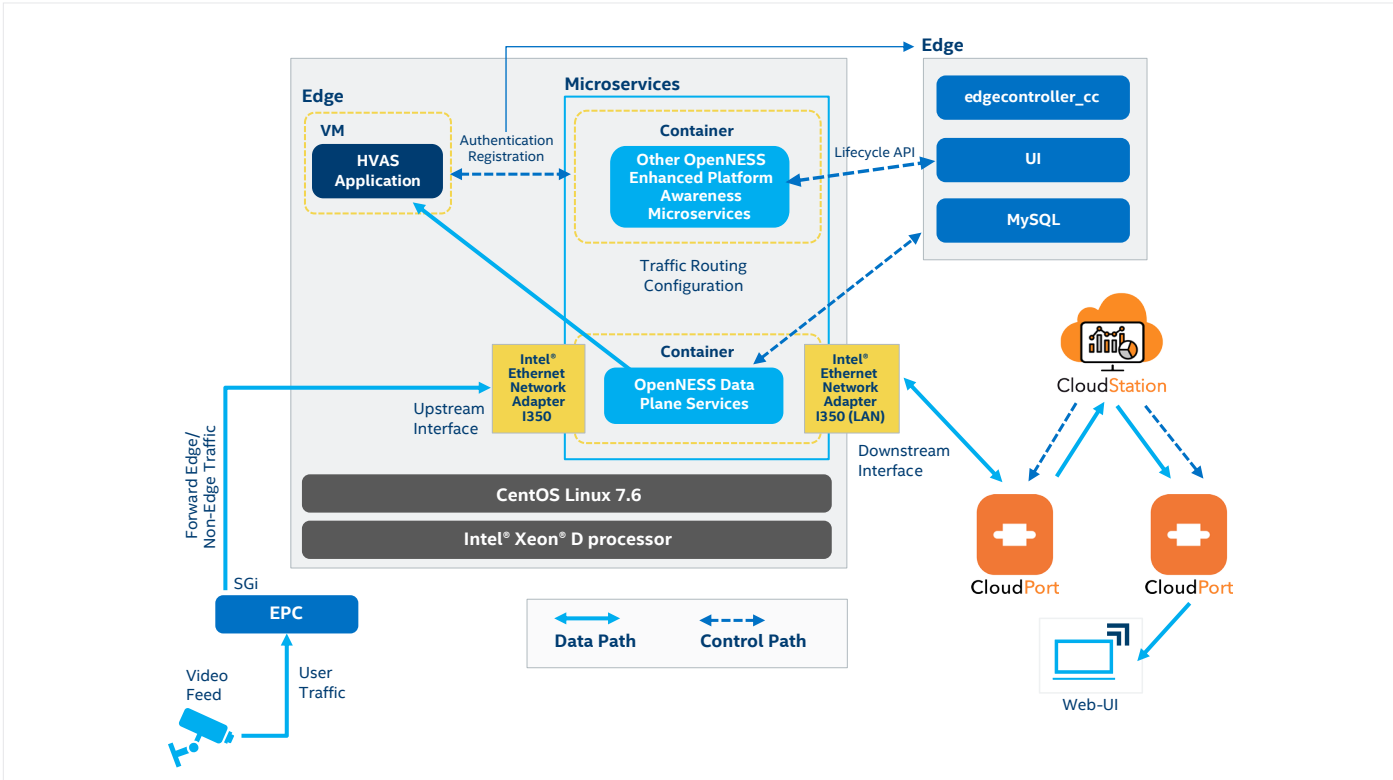

Figure 4. OpenNESS integration with HVAS application.

Data and Control Flows

HVAS integration with OpenNESS has two flows as shown in Figure 4.

1. **Data flow:** As per the 3GPP standards, the video feed terminates on a base station (eNodeB for 4G or a gNodeB for 5G) and then connects to the OpenNESS edge node via an SGi interface (4G LTE) or an N6 interface (5G). In the HVAS implementation, video feeds reach the HVAS application on the edge node through the upstream interface. This video traffic is intercepted by the data plane of the edge node data plane services. On the OpenNESS edge controller, traffic policies are configured such that traffic steering is applied to either redirect the traffic to edge applications for further analysis or to pass the packets through to the downstream interface. The downstream interface is connected to the SD-WAN edge device. Normal traffic is forwarded via the SD-WAN edge devices with no prioritization. When a security incident has occurred, a notification message is triggered. The SD-WAN controller is configured to act on this notification message and push a configuration prioritizing this traffic to all SD-WAN edge devices. The monitoring station is connected to a remote SD-WAN edge device.
2. **Control flow:** The application is authenticated with the edge node appliance (which includes ELA, EVA, EDA and EAA) and the edge controller. Once authentication is successful, the application is registered. In lifecycle management, the ELA communicates with the edge controller to control the status of the application (start, stop, delete, etc.). Traffic routing configuration is defined on the controller. Based on the traffic, policy routing decisions are made: incoming video camera traffic is redirected to the HVAS and other traffic is sent directly to the SD-WAN edge device without any HVAS processing.

Figure 5 shows how HVAS is deployed on the edge node. Once the interfaces are declared as user plane interfaces, the HVAS application VM is deployed on the edge node using the controller UI. Once the HVAS VM status is changed to running state, the source filter is applied on the HVAS VM to redirect traffic to the HVAS application. This processing is done entirely at the edge and any detected security violations are sent to remote monitoring personnel.

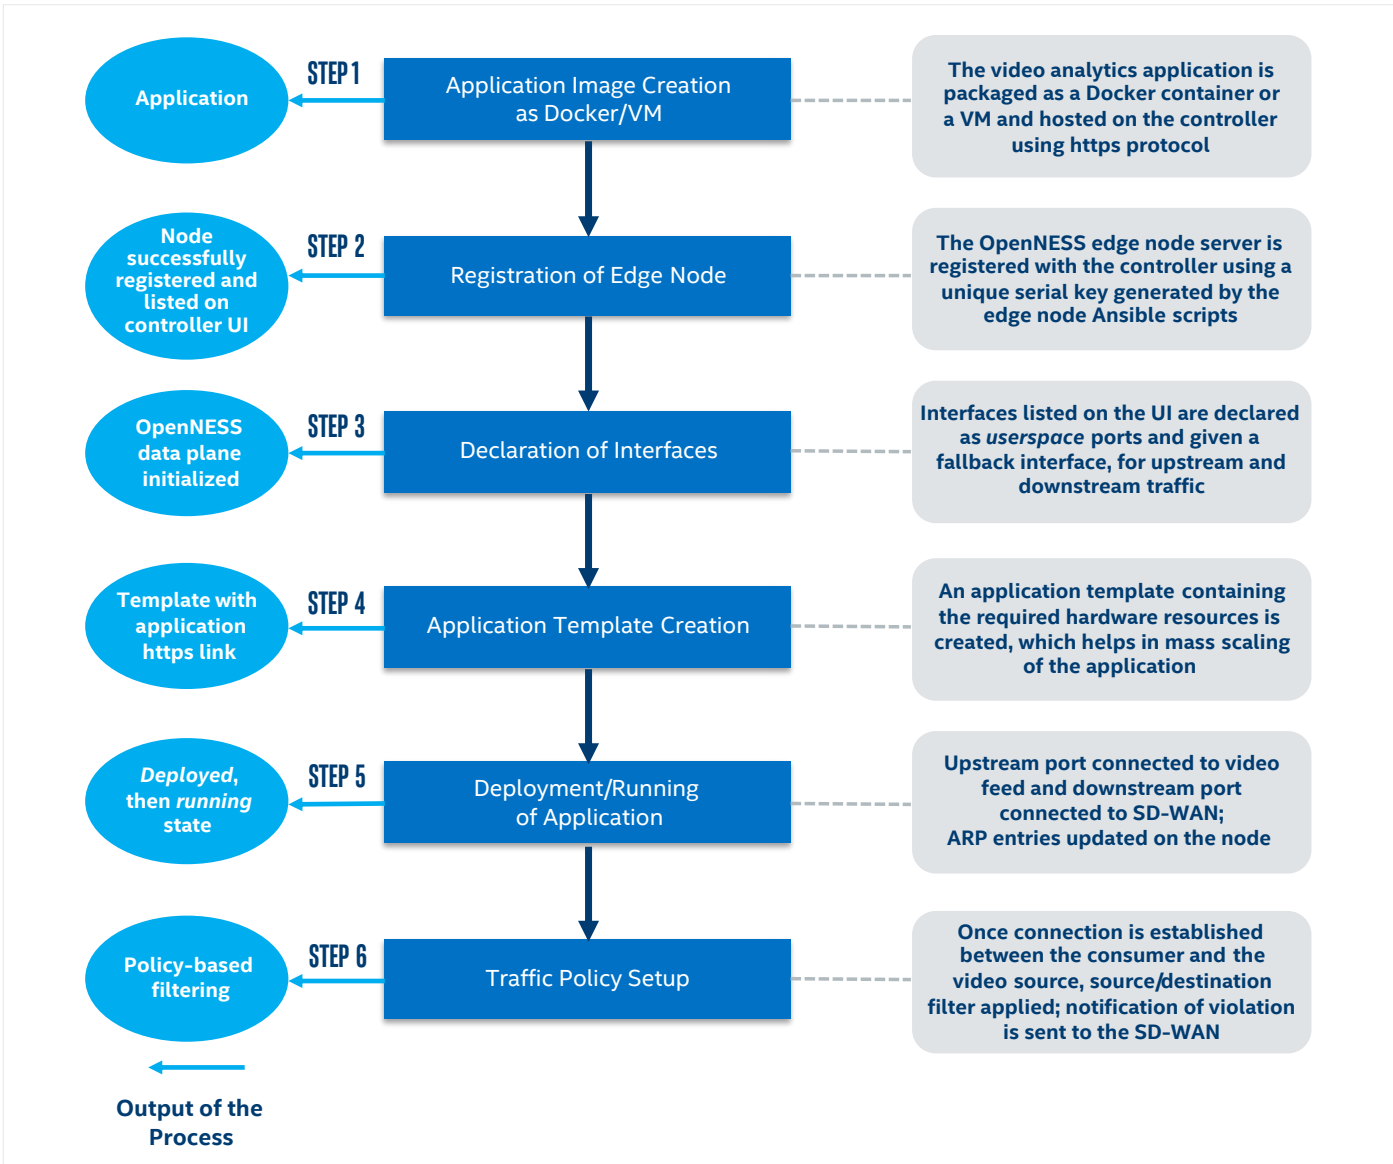

Figure 5. Flow diagram for application deployment on edge.

Conclusion

Wipro has developed the HVAS solution leveraging edge analytics to utilize camera networks for real-time alerts in a way that reduces backhaul network bandwidth utilization, reduces latency, and improves scalability so that the learning model obtained from one node can be propagated across all instances of the service. By using OpenNESS, computation offloading at the on-premises edge is possible, allowing video processing locally to minimize backhaul traffic. Orchestration can be leveraged to replicate the HVAS solution across multiple edge instances based on need.

This HVAS application enables new possibilities of providing better protection to ATM machines, where the HVAS can be trained to detect fire or water or anything that is not normal and send out alert notifications instead of full video traffic. In the case of using HVAS for police stations, the HVAS can be trained to detect guns or knives and provide warning, and it can be a deterrent. HVAS delivers the value of helping prevent unfortunate, controllable incidents that may lead to loss of life or tragedy.

About Wipro

Wipro Limited (NYSE: WIT, BSE: 507685, NSE: WIPRO) is a global information technology, consulting, and business process services company. It harnesses the power of cognitive computing, hyper-automation, robotics, cloud, analytics, and emerging technologies to help its clients adapt to the digital world and make them successful. A company recognized globally for its comprehensive portfolio of services, strong commitment to sustainability and good corporate citizenship, it has have over 160,000 dedicated employees serving clients across six continents. Together, its employees and clients discover ideas and connect the dots to build a better and a bold new future.

About Intel® Network Builders

Intel Network Builders is an ecosystem of infrastructure, software, and technology vendors coming together with communications service providers and end users to accelerate the adoption of solutions based on network functions virtualization (NFV) and software defined networking (SDN) in telecommunications and data center networks. The Network Edge Ecosystem is an initiative gathering ecosystem partners with a focus on accelerating network edge solutions. As an integral part of the broader Intel Network Builders program, this initiative aims to facilitate partners' access to tested and optimized solutions for network edge and cloud environments. Learn more at http://networkbuilders.intel.com/networkedgeecosystem

TABLE OF ABBREVIATIONS

3GPP	3rd Generation Partnership Project
API	Application Programming Interface
CNN	Convolutional Neural Network
DPDK	Data Plane Development Kit
EAA	Edge Authentication Agent
EDA	Edge Dataplane Agent
ELA	Edge Lifecycle Agent
EPC	Evolved Packet Core
EVA	Edge Virtualization Agent
HVA	High Value Asset
HVAS	High Value Asset Surveillance
LAN	Local Area Network
MEC	Multi-Access Edge Computing
NFV	Network Functions Virtualization
ONAP	Open Networking Automation Platform
OpenNESS	Open Network Edge Services Software
QoS	Quality of Service
SD-WAN	Software Defined WAN
SDN	Software Defined Networking
SLA	Service-Level Agreement
VAE	Video Analytics Engine
WAN	Wide Area Network