WHITE PAPER

Service Providers
NFV

Simplify Application Development for the Network Edge

Hosting applications at the network edge enables innovative low latency use cases. However, ingesting data has traditionally required an understanding of how the network works. An Intel® SDK reduces the complexity

Table of Contents

Executive Summary

There is an opportunity for Communications Service Providers (CoSPs) to enhance the user experience and create new revenue streams by deploying applications at the edge of the network, significantly cutting latency. Multi-access Edge Computing (MEC) initiatives work in tandem with Network Functions Virtualization (NFV), using compute capacity on general purpose processors to add intelligence to the network.

For developers, hosting applications in the network introduces complexities they do not encounter in the cloud, in particular with regard to ingesting data. The Intel® Network Edge Virtualization Software Development Kit (Intel® NEV SDK) provides APIs and libraries that abstract away the complexity, so that developers can build applications without needing an understanding of the underlying network architecture and function.

Solution Benefits

- **Deliver new user experiences.** The Intel® Network Edge Virtualization Software Development Kit (Intel® NEV SDK) makes it easier to develop applications for new use cases now.
- **Create a foundation for 5G.** 5G, supported by edge computing, has the potential to enable new applications including autonomous vehicles, augmented reality and virtual reality. Today's edge applications can be enhanced with the extra bandwidth of 5G in the future.
- **Abstract away complexity.** Create applications for the edge without needing an understanding of the underlying network topology or function.
- **Integrate with Amazon Web Services* (AWS*).** Using the provided lambdas, the Intel NEV SDK enables developers to ingest data from the network for use in cloud applications hosted at the edge.

Tackling Complexity at the Network Edge

While 5G dramatically increases the bandwidth available for innovative applications, the latency associated with backhaul across the network will remain a constraint. MEC initiatives work in tandem with NFV initiatives, to host applications

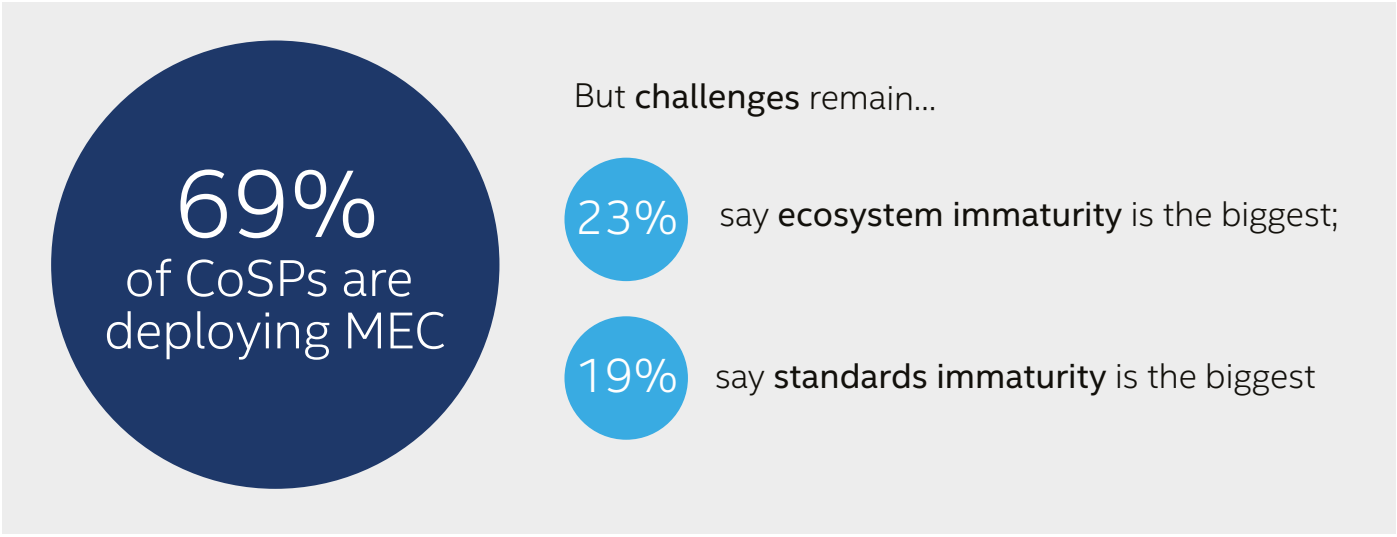

Figure 1. Ecosystem and standards immaturity remain challenges in Multi-access Edge Computing (MEC), even as many Communications Service Providers (CoSPs) make plans for deployment

on general purpose hardware, closer to the edge of the network. This cuts the latency that end users experience and can also reduce the CoSP's costs by minimizing backhaul. Technologies such as autonomous vehicles, augmented reality and virtual reality will depend on edge computing for rapid response. Additionally, applications such as smart factories, video transcoding, and Content Delivery Networks (CDNs) can benefit from local solutions that can reside at the edge of the network.

The term "network edge" masks a lot of complexity, and is used to refer to locations that sit between the user devices and the core CoSP network. It could include small cells, aggregation sites, macro cells, central offices, customer premises equipment and the Radio Access Network (RAN). Edge computing locations range from edge compute nodes hosted at millions of on-premise sites; network or other intermediary data hubs; and regional data centers or co-location sites. Virtualized servers can provide a common infrastructure across these network locations, with compute capacity available for running applications alongside Virtual Network Functions (VNFs).

From the application developer's point of view, creating software for the network can be more complex than creating software for the cloud. If the edge computing resources sit close to the base station, the data might need to be untunneled. Any IPSec stream will need to be terminated, and then the data must be extracted. Another option to insert the data plane closer to the edge would be to implement a Control and User Plane Separation (CUPS) deployment model, which allows the control plane to be largely unchanged. In an Internet of Things (IoT) environment, data might be coming in using the Message Queuing Telemetry Transport (MQTT) protocol. Deeper in the network, there might be a Gateway Internet Local Area Network (SGi-LAN) interface to contend with. There is a need for a consistent and standardized way to access the data plane, that abstracts away the complexity, in order to make the network a desirable development platform for cloud-based application developers.

In the next two years, 69 percent of CoSPs will be deploying MEC[1]. However, challenges remain. The ecosystem is considered immature by 23 percent of CoSPs, and 19 percent of them say standards are immature (see Figure 1). There is a need for open APIs that can abstract away the complexity, and enable CoSPs to more easily collaborate with innovative software developers.

Standard APIs for Network Access

To help accelerate the maturity of the ecosystem and standards for edge computing, Intel has created the Intel NEV SDK.

The Intel NEV SDK helps CoSPs and application developers to create MEC applications. It provides open Application Programming Interfaces (APIs) and libraries for easily ingesting data at the network edge and abstracts away the complexity of hosting applications at different points in the network. As a result, developers can create applications for the network without needing to understand the underlying mobile network protocols, and those applications will work wherever they are hosted in the network. Using their existing software engineering investment and expertise, developers can create applications for deployment in the network, opening up new revenue and customer service opportunities.

Developers who use Amazon Web Services* (AWS*) can use the Intel NEV SDK together with AWS Greengrass* to host applications close to the edge. Their applications can be enhanced with network access using AWS lambdas provided as part of the SDK. The lambdas included with the NEV SDK bind the Amazon Greengrass Core to the underlying

MEC platform. Customer specific lambdas can run on the Greengrass Core without needing to understand the underlying networking protocols.

The SDK enables new services to be created, such as location-aware applications, and creates opportunities for CoSPs to increase revenue and reduce churn in consumer and enterprise markets.

Use Case: Facial Recognition in Retail

Intel, Foxconn Technology Group and Asia Pacific Telecom have demonstrated facial recognition applications in retail, based on the Intel NEV SDK (see Figure 2). The proof-of-concept showed how shoppers could use a facial scan for applications including cashless payment using facial recognition. The technology can also be used to enable access control.

In the demonstration, wireless cameras connected to the Internet using small cell radio access nodes and transmitted the shopper's image to the local MEC platform, where it was processed. Hosting the application in the MEC platform helped to cut latency and improve the customer experience. According to Foxconn Technology Group and Asia Pacific Telecom, the payment authentication process can be completed within 0.03 seconds, making for a premium customer experience. Payment based on facial recognition can also help to reduce credit card fraud while reducing the risk of personal information leakage.

Solution Architecture: APIs for the Data Plane

The Intel NEV SDK runs on the Intel® Xeon® processor, widely used by CoSPs for existing VNFs. An application is hosted with the SDK in a virtual machine (VM) or container, on a server that also hosts VNFs in other virtual machines. Authenticated applications can communicate with each other and with VNFs over vSwitches. The server can also include Field Programmable Gate Arrays (FPGAs) or other accelerators.

AWS provides Amazon Greengrass to enable IoT applications to run seamlessly across the cloud and local devices. The Amazon Greengrass Core can be hosted on a virtual machine or container on the edge server, enabling AWS Lambda functions to run there. The Intel NEV SDK provides AWS Lambda functions to enable low latency data plane access. As a result, cloud applications can be more easily migrated to the network and extended with network capabilities using a developer's existing tools and skills.

The Intel NEV SDK can also work together with the OpenVINO™ toolkit for computer vision, which can be hosted in a separate VM or container on the same server. For example, an application could use the OpenVINO toolkit for object detection and classification, and the Intel NEV SDK to receive security camera traffic. Working together with the Intel® Media SDK, the Intel NEV SDK can be used for video transcoding applications.

Figure 2. Intel, Foxconn Technology Group and Asia Pacific Telecom demonstrated facial recognition systems for retail, based on the Intel® Network Edge Virtualization Software Development Kit (Intel® NEV SDK), in Taiwan

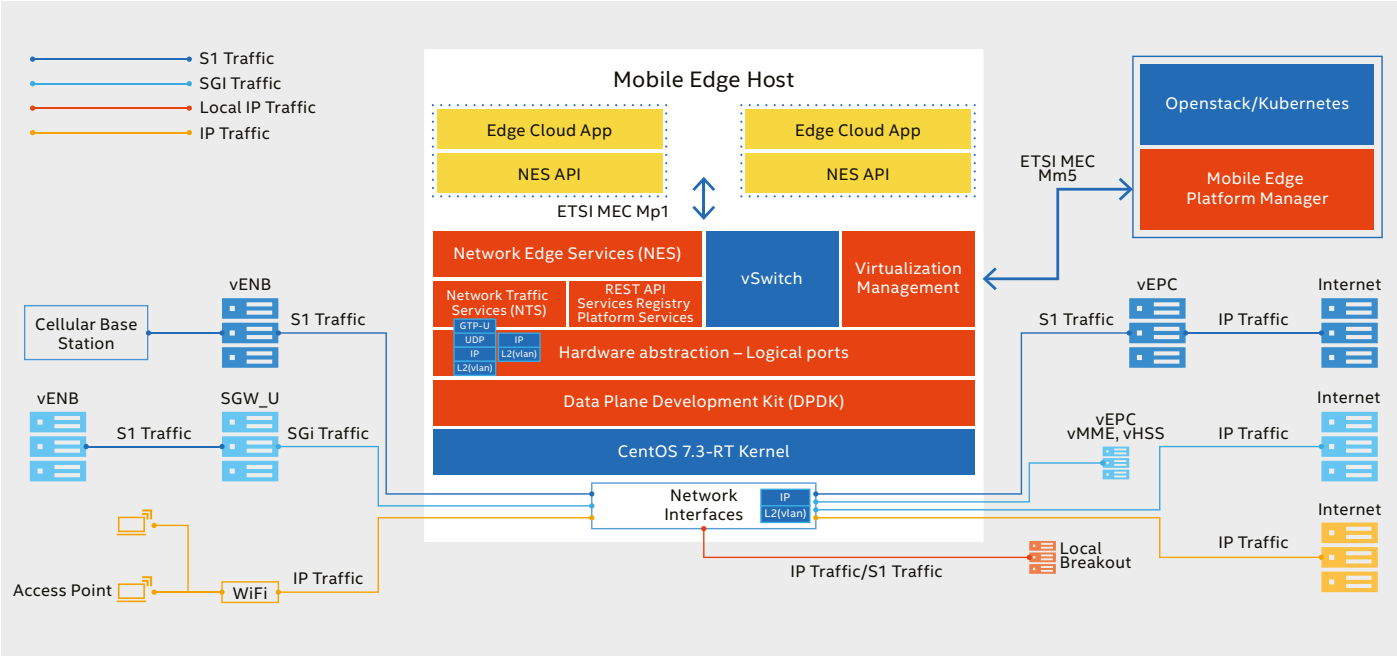

Figure 3. The architecture for hosting applications at the edge of the network using the Intel® Network Edge Virtualization Software Development Kit (Intel® NEV SDK)

Applications can be orchestrated in the infrastructure using OpenStack* and Kubernetes* (see Figure 3). Intel has validated multiple best-known configurations, including one that uses a dual-socket server based on the Intel® Xeon® Gold 6148 processor in a 2U rackmount server enclosure, validated with the Wind River Titanium* Cloud NFV and virtualization platform. The Intel® Ethernet Converged Network Adapter X710-DA4 and the Intel® Ethernet Network Adapter XXV710-DA2 are used for networking, together with the Intel® Ethernet Network Connection OCP X557-T2. The server is fitted with the Intel® QuickAssist Adapter Device.

The Intel NEV SDK can be deployed across the network, including in small cells, aggregation sites, macro cells and the C-RAN.

Conclusion

Using the Intel NEV SDK, CoSPs can simplify the development of enterprise applications for the network edge. By enabling application developers to bring their software to the network, CoSPs can create new revenue opportunities and prepare the network to deliver the processing-intensive applications that 5G will enable.

Intel will be participating in industry initiatives to help increase openness at the edge of the network, and the Intel NEV SDK can make a valuable contribution.

Learn More

- Intel® NEV SDK
- Facial-Recognition Technology at the Cutting Edge of Retail, High-Speed Authentication
- Intel® Network Builders
- Intel® Xeon® processors
- Intel® Xeon® Gold 6148 processor
- Intel® Ethernet Converged Network Adapter X710-DA4
- Intel® Ethernet Network Adapter XXV710-DA2
- Intel® Ethernet Network Connection OCP X557-T2
- Intel® QuickAssist Adapter Device
- OpenVINO™ toolkit
- Intel® Media SDK

Find the solution that is right for your organization. Contact your Intel representative or visit **https://www.intel.com/content/www/us/en/communications/network-transformation.html**

Solution Provided By:

[1] Transforming the Edge: The Rise of MEC, A Heavy Reading white paper produced for Intel.

0918/CAT/GC/PDF 338156-001EN