Solution Brief

Intel® Network Builders
Telecommunications

Building On-Demand Edge Clouds from Mixed Infrastructure

The MobiledgeX Edge-Cloud platform provides end-to-end control and management of network edge workloads across multi-cloud edge deployments. This universal orchestration and automation solution enables communications services providers to gain a unified view and full control of security, network addressing scheme, IP keys, privacy, and data governance for full end-to-end control of their network edge infrastructures and workloads. MobiledgeX collaboration with Intel integrates the latest Intel reference architecture and chipset accelerators to optimize performance, decrease latency, and increase security and efficiency of edge networks.

Communication service providers (CSPs) building out 5G initiatives have dramatically accelerated their adoption of public cloud infrastructure, a trend that is likely to continue for the foreseeable future. At the same time, the prospect of massive increases in traffic rates and data volumes, coupled with the need to meet strict latency requirements, are pushing compute tasks out to the network edge, leading to accelerated CSP build-outs of multi-access edge computing (MEC) infrastructure. Distributed workloads at the network edge reduce delay, bandwidth costs, and traffic congestion by reducing backhaul traffic to the network core. They also improve support for latency-sensitive workloads, such as those that involve real-time collaboration and control.

To help realize the competitive potential of these technology shifts, The MobiledgeX Edge-Cloud platform provides a universal orchestration platform and unified control plane to manage and control edge compute workloads across any cloud infrastructure. The MobiledgeX Edge-Cloud platform enables CSPs to create edge clouds from any combination of their own edge infrastructure and public cloud resources, as illustrated in Figure 1.

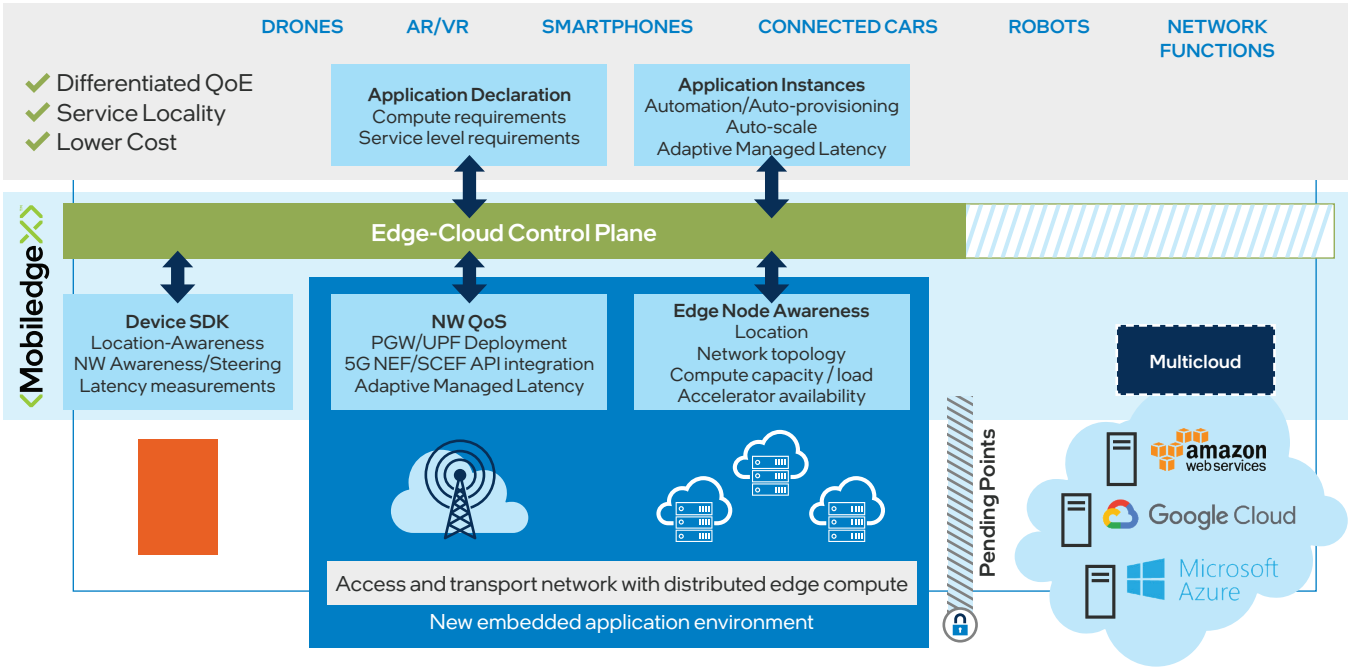

Figure 1. MobiledgeX management and control of federated edge clouds.

This approach allows CSPs to retain full control while being able to partner with multi-cloud and hybrid cloud infrastructure providers. Edge cloud resources can be used for internal workload functions, or they can be offered to customers to deploy and execute their own workloads across multi-cloud edge environments tailored to their individual needs. The MobiledgeX Edge-Cloud platform manages applications and workloads across heterogeneous topologies while protecting data privacy and security. Application developers onboard their services globally and declare their own defined service levels. Then the MobiledgeX Edge-Cloud solution dynamically orchestrates edge services to the appropriate developer selected to enable network elements and edge cloud infrastructure based on the client's request. By leveraging signaling from the client side, the MobiledgeX SDK, and infrastructure telemetry, the application can deliver the desired quality of experience (QoE).

Enablement work with Intel is developing proof points and use cases for deployments of the MobiledgeX Edge-Cloud solution with 3rd Generation Intel® Xeon® Scalable processors. These implementation proofs of concept add to the ecosystem of highly optimized managed edge solutions and offer the following benefits to CSPs:

- **Flexibility**. CSPs can deploy federated edge clouds based on their own edge cloud resources, public cloud instances, or any combination, optimizing use of owned infrastructure with unlimited scalability.
- **Control**. The MobiledgeX Edge-Cloud solution provides cohesive, unified oversight and control over aggregated edge resources, applications, and workloads, without integration challenges across infrastructure providers.
- **Profitability**. New revenue streams can be accelerated to harness and monetize new access to network services available on the network edge. CSPs can now participate across the entire edge computing value chain instead of just providing bandwidth.

This enablement builds on joint work contributed by Intel and MobiledgeX according to the GSMA Operator Platform Edge Requirements specifications as participants of the GSMA Operator Platform Group.

Use Case: Device Offloading for Foot Traffic Analysis

Device offloading consists of shifting compute-intensive workloads such as video analytics from endpoints such as cameras to edge clouds for processing. Data science and AI/ML solution developer Kibernetika deploys its Heatmap Foot Traffic Analysis solution using MobiledgeX Edge-Cloud on 3rd Generation Intel Xeon Scalable processors. The solution connects to multiple cameras in a retail space and can use any combination of cloud infrastructure: on-prem at the customer location, on CSP-operated edge infrastructure, or on public clouds.

The Kibernetika solution analyzes customer foot-traffic patterns from the camera video streams, combined with transaction data from point-of-sale (PoS) devices, to quantify in-store customer behavior in depth. It uses these analytics to generate visualizations use such as heatmaps and track maps that deliver insights to retailers. Store operators can apply that understanding to help guide store layouts and product placement for optimized sales. Analytics from the solution have other retail applications as well, such as managing customer waiting time, detecting shoplifters, and monitoring product levels on shelves.

The video analytics required to generate audience analytics heatmaps are particularly compute intensive, and scalability of the solution is vital in retail environments that may include many cameras at each of many sites. Kibernetika optimizes its Machine Teaching technology for Intel architecture by adopting toolkits developed alongside Intel architecture. Intel® Distribution of OpenVINO™ toolkit provides capabilities to build, optimize, and run deep-learning inference models on platforms based on Intel® processors and accelerators. The toolkit enables deep learning inference based on convolutional neural networks (CNNs) at the network edge, accelerating time to market using a library of readily adopted computer functions and pre-optimized kernels.

3rd Generation Intel® Xeon® Scalable Processors

- **Flexibility from the edge to the cloud**, with a balanced architecture, built-in acceleration, and hardware-based security.
- **Part of a complete set of network technology from Intel**, including accelerators, Ethernet adapters, Intel® Optane™ persistent memory, FlexRAN, Open Visual Cloud, and Intel® Smart Edge.
- **Engineered for 5G network functions**, targeting low latency, high throughput, deterministic performance, and high performance per watt.
- **Enhanced built-in crypto-acceleration** including Intel® QuickAssist Technology (Intel® QAT) to increase performance of encryption-intensive workloads.
- **Hardware-based security** using Intel® Software Guard Extensions (Intel® SGX),[1] enhanced crypto processing acceleration,[1] and Intel® Total Memory Encryption.[1]

A Kibernetika test team quantified that value, finding that the test server could support 144 camera streams for the audience analytics heatmap workload based on the Intel Distribution of OpenVINO toolkit, compared to 96 streams for the machine-learning workload based on the PyTorch framework, as shown in Figure 2.[2]

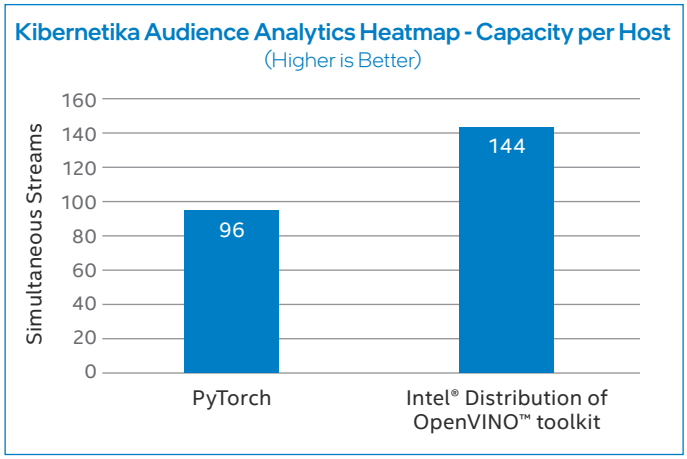

Figure 2. Capacity (simultaneous streams per server) for Kibernetika heatmap generation.[2]

Cloud-Native 5G User Plane Function (UPF)

Deutsche Telekom, MobiledgeX, and Mavenir collaborated on a proof-of-concept design for a cloud-native UPF for edge applications on CSP-operated or private 5G networks as well as in public clouds. The MobiledgeX Edge-Cloud 3.0 platform dynamically places UPF workloads on Intel® architecture-based edge infrastructure that is tuned to workload-specific requirements for proximity and performance.

More information ›

Conclusion

The ability to define federated edge clouds based on any combination of owned edge assets and public cloud infrastructure is an ongoing monetization opportunity for CSPs. The MobiledgeX Edge-Cloud solution is a flexible, universal orchestration platform and control plane for this topology, optimized for performance on Intel architecture. This combination allows CSPs to offer edge cloud services with optimized TCO and service levels while retaining full control. On-demand edge clouds have potential to deliver new revenue while improving utilization of capital investments.

More Information

Intel® Network Builders: **networkbuilders.intel.com**

Intel Distribution of OpenVINO Toolkit:
software.intel.com/content/www/us/en/develop/tools/openvino-toolkit.html

MobiledgeX Edge-Cloud Platform: **mobiledgex.com/product/**

Kibernetika and Intel AI Builders: **builders.intel.com/ai/membership/kibernetika**

Solution provided by:

[1] This technology is not supported when using Intel® Optane™ persistent memory.

[2] Testing done by Kibernetika Inc. on 26 June 2021.

System under test: One node, 2x Intel® Xeon® Gold 6230N processors (20 cores/40 threads, 2.30 GHz, 27.5 MB cache); Intel® Hyper-Threading Technology enabled; Intel® Turbo Boost Technology disabled; 384 GB RAM; CentOS 3.10.0; Heatmap using Pytorch and Intel Distribution of OpenVINO™ toolkit models; GCC 4.8.5.

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for configuration details. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a nonexclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

1021/DO/MESH/346429-001US