Solution Brief

Intel® Network Builders
Telecommunications

Extending Cloud Services to the Network Edge

NxtGen SpeedCloud™ Edge brings cloud-based compute power to the edge, running workloads closer to users for a superior experience. Leveraging Intel® Smart Edge Open and other open source components, the solution is optimized for Intel® architecture to help enterprises improve performance and scalability while reducing their operating costs.

NxtGen Infinite Datacenter

Massive increases in data volumes strain the viability of traditional approaches to centrally aggregating that data for processing at the network core. Bandwidth requirements to backhaul massive amounts of data from fast-proliferating endpoint sources can quickly become cost-prohibitive, limiting potential value. The inherent delay associated with transporting data back to a central area for processing also interferes with latency-sensitive usages such as communications, media, and real-time analytics.

To address these challenges, solution architects increasingly push processing out to the network edge. Working on data close to its source reduces bandwidth costs and latency, enabling edge-based distributed services to mitigate traffic congestion and safeguard application performance. The full benefit of such edge deployment depends on the ability to efficiently and consistently deploy and manage infrastructure and applications across the infrastructure while optimizing performance, responsiveness, and security.

Solution Architecture to Enable the Edge Computing Paradigm

NxtGen SpeedCloud Edge simplifies the deployment and management of container-based applications and workloads at the network edge. The platform provides a distributed, centrally managed Kubernetes infrastructure, Intel® Smart Edge Open, the Intel® Distribution of OpenVINO™ toolkit, and OpenNebula.

SpeedCloud Edge is part of a broader set of offerings from NxtGen that also includes SpeedCloud Core. Together, these components form the SpeedCloud public cloud computing platform. The solution abstracts away the complexity of provisioning and maintaining cloud infrastructure, so IT organizations can increase their effectiveness by focusing on more value-added activities. The SpeedCloud platform as a whole enables management of the cloud from edge to core as a single entity, as illustrated in Figure 1.

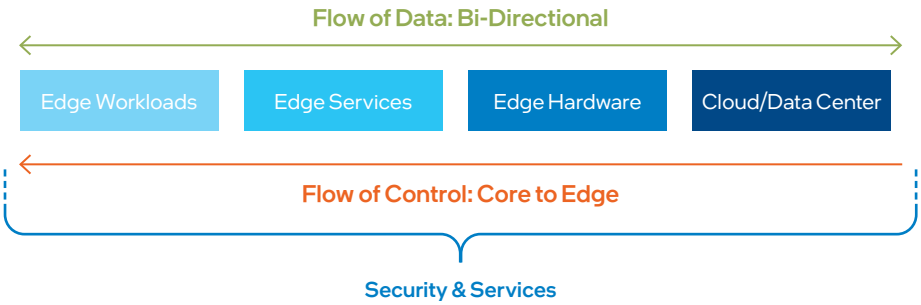

Figure 1. High-level NxtGen SpeedCloud™ solution architecture.

A hub-and-spoke topology enables centralized management and control of the distributed SpeedCloud Edge environment, as illustrated in Figure 2. The spokes are compute-enabled edge nodes hosted by a global network of data center partners, enabling SpeedCloud Edge to deploy workloads geographically. This approach delivers high performance and low latency wherever it is needed, while also offering regionalized services and meeting local standards and regulatory requirements. Centralized orchestration provisions and manages edge nodes and automates the deployment and lifecycles of containerized workloads.

The commitment to open source software components within the SpeedCloud technology stack helps foster extensibility and interoperation with other solutions, interfacing with the broader cloud ecosystem and helping ensure the platform's future-readiness. Common use cases and capabilities realized by customers using SpeedCloud Edge include the following:

- **Holistic, centrally managed container infrastructures** based on Kubernetes clusters that enables customers to provision web and application workloads in a self-serve modality with automatic deployment to the cloud.
- **Edge-based ingestion and processing of IP video feeds** for usages such as transcoding and analytics, reducing network bandwidth requirements and costs by avoiding backhaul of massive data sets to the network core.
- **Accelerated content delivery networks (CDNs),** with robust processing at the edge for data aggregation and delivery of rich cloud gaming and media content for improved customer experiences.
- **Edge-based microservices on Kubernetes clusters,** providing services such as differential content delivery based on geography, such as enabling regional offerings by eCommerce providers.

Hardware and Software Stack

Ongoing collaboration between NxtGen and Intel includes enablement work to help ensure high performance and reliability from applications distributed at the network edge using SpeedCloud Edge and Intel platforms.

Open Source Software Stack, Optimized by Intel

The open source software ecosystem provides the foundations on which the SpeedCloud Edge platform is built and enables optimized performance on Intel architecture:

- **Intel® Smart Edge Open (formerly known as OpenNESS)** is a royalty-free edge computing software toolkit that enables highly optimized and performant edge platforms to on-board and manage applications and network functions with cloud-like agility across any type of network. Intel Smart Edge Open includes functional building blocks and experience kits to provide flexibility and stability in edge solutions according to customers' unique needs.
- **Intel® Distribution of OpenVINO™ toolkit** wstreamlines the process of building, optimizing, and running deep-learning inference models based on convolutional neural networks (CNNs) at the edge, including a library of pre-optimized kernels and computer functions.

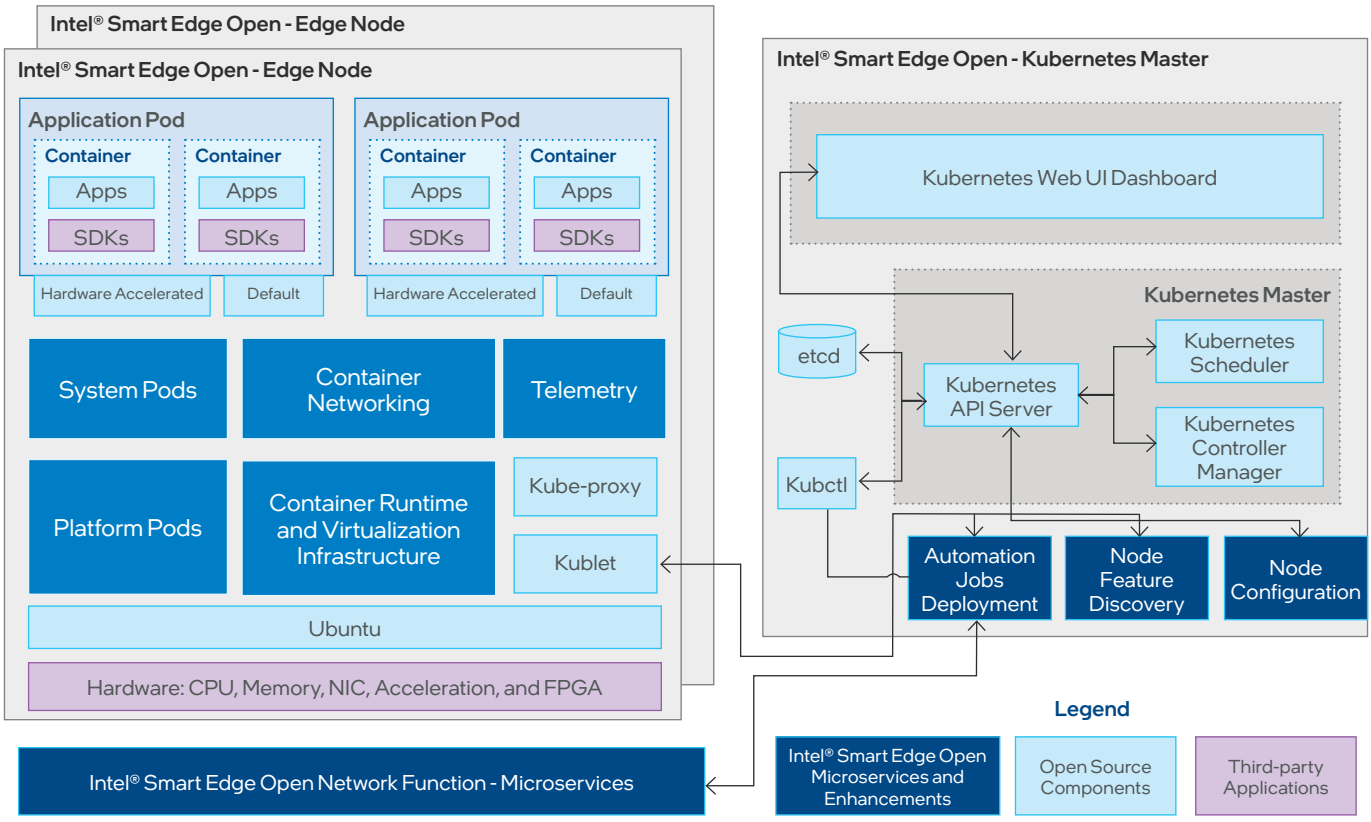

Figure 2. NxtGen SpeedCloud™ Edge technology stack.

Management from the Network Core to the Edge

OpenNebula provides a single point of management for cloud infrastructure, from the core to the edge, with automated resource provisioning, workload portability across cloud and edge providers, and interoperability across providers to create an aggregated service platform.

High-Performance Integrated Hardware Platform

NxtGen SpeedCloud Edge is built to deliver optimized performance on Intel architecture-based platforms. Key components of the hardware stack used for SpeedCloud Edge deployments include the following:

- **Intel® Xeon® Scalable processors** offer flexibility for edge workloads, with options at a variety of feature levels and core counts. The processors provide high per-core performance and accelerate machine learning functions using Intel® Deep Learning Boost with Vector Neural Network Instructions (VNNI).
- **Intel® Optane™ persistent memory** redefines system memory and storage architectures with a large, affordable persistent tier that enables operators to scale system memory using modules up to 512 GB that are socket-compatible with conventional DRAM.
- **Intel® Movidius™ Myriad™ X vision processing unit (VPU)** accelerates machine learning-driven visual workloads using a dedicated, on-chip Neural Compute Engine. The VPU processing pipeline is built to be programmed using the Intel Distribution of OpenVINO toolkit.

Conclusion

Using NxtGen SpeedCloud Edge, organizations can streamline the deployment of applications and workloads to the network edge using a robust, cloud-native model based on open source software components and optimized for Intel infrastructure. Placing compute and storage at the edge reduces costs associated with wide-area network bandwidth consumption and enables latency-sensitive usages based on highly distributed data sources.

Container infrastructure based on SpeedCloud Edge helps streamline network transformation, adding global scale on-demand to edge computing environments. SpeedCloud Edge can also be a key enabler for container-as-a-service (CaaS), which lets development teams deploy containerized software without concerning themselves with lower-level infrastructural details. These models and others offer a robust on-ramp for businesses to profit from the symbiotic relationship between cloud computing and processing at the network edge.

More Information

Intel® Network Builders: networkbuilders.intel.com

NxtGen SpeedCloud™: speedcloud.co.in

NxtGen Multiverse: multiverse.nxtgen.com

NxtGen Blog – Edge Computing, Simplified: nxtgen.com/edgecomputing-simplified.php

Solution provided by:

intel + NxtGen Infinite Datacenter

1221/DO/MESH/346431-001US