

# Next Generation Central Office - Converged Edge Architecture



## stc Digital Strategy and Objectives



stc is a pioneer and digital champion in the Middle East and North Africa (MENA) region and is the leading telecommunication service provider in the country. stc's focus on innovation and evolution helps to insure they remain at the forefront of change.

Technology transformation is a major enabler of stc's overarching strategy, which is to be a world-class digital leader that provides innovative services and platforms that enable the digital revolution for the Kingdom of Saudi Arabia and the MENA region. Edge infrastructure strategy is pivotal to the success of stc's digital transformation journey. stc channeled their focus to leverage its access footprint across the Kingdom, and up-leveled its Information and Communications Technology (ICT) service offerings to meet the current demand. The transformation of stc to a digital service provider (DSP) aligns with the goals of their customers, including large enterprises and government, and supports the National Transformation Program and Saudi Vision 2030.

True to its strategy and vision, stc used technology as an enabler to open new horizons for market penetration. Bearing in mind the cost and financial challenges facing most existing communications service providers (CSPs) around the globe, stc paved the way for the industry to enable FA (fixed access) and MA (mobile access) virtual functions convergence on a shared infrastructure. This provides DSPs the ability to dynamically allocate their infrastructure resources for fixed and mobile workloads based on business needs. This enables stc to provide a unified digital service experience across fixed and mobile access, allowing for a true fixed and mobile convergence from both an infrastructure and service offering perspective. To demonstrate the business value realization from this technology adoption, stc collaborated with Intel, Nearby Computing, VMware and Cisco on a proof-of-concept project, based on Intel's Converged Edge Reference architecture.

## Introduction: Transforming the Central Office and Converged Access

Unlike before, today's broadband-dominated mobile traffic is becoming increasingly disengaged from the revenue growth. Legacy services such as voice, data and video services are diminishing in value. Subscribers pay less per bit for these services, resulting in flat or diminishing revenue growth model, and a much lower level-off in the costs required to support this traffic growth.

Globally, CoSPs have witnessed the costs of traditional communications equipment drop, while traffic growth continues unabated. Hence, CoSPs must now explore alternative architectural approaches and business models toward reducing costs while building out new flexible software-based architectures, that allow them to quickly and efficiently offer and grow new replacement service and revenue streams.

The central office location is the ideal deployment site for this type of approach. The fiber rich access to the location produces the ideal environment for these new, higher bandwidth, higher value, and low latency edge service deployments.

Moving forward, the level of service and edge distribution will vary depending on the services offered by stc. We expect to see the mobile infrastructure distributed to the central office location in order to serve the massive and sustained 5G traffic CAGR, as well as offering fixed mobile services. For very low latency edge

### Table of Contents

- stc Digital Strategy and Objectives... 1
- Introduction: Transforming the Central Office and Converged Access ..... 1
- Objective of the stc Proof of Concept..... 4
- Intel's Approach to the Converged Edge Architecture..... 6
- The Nearby One Solution: Orchestration for the Modular Edge. . 6
- Conclusion and Next Steps ..... 7

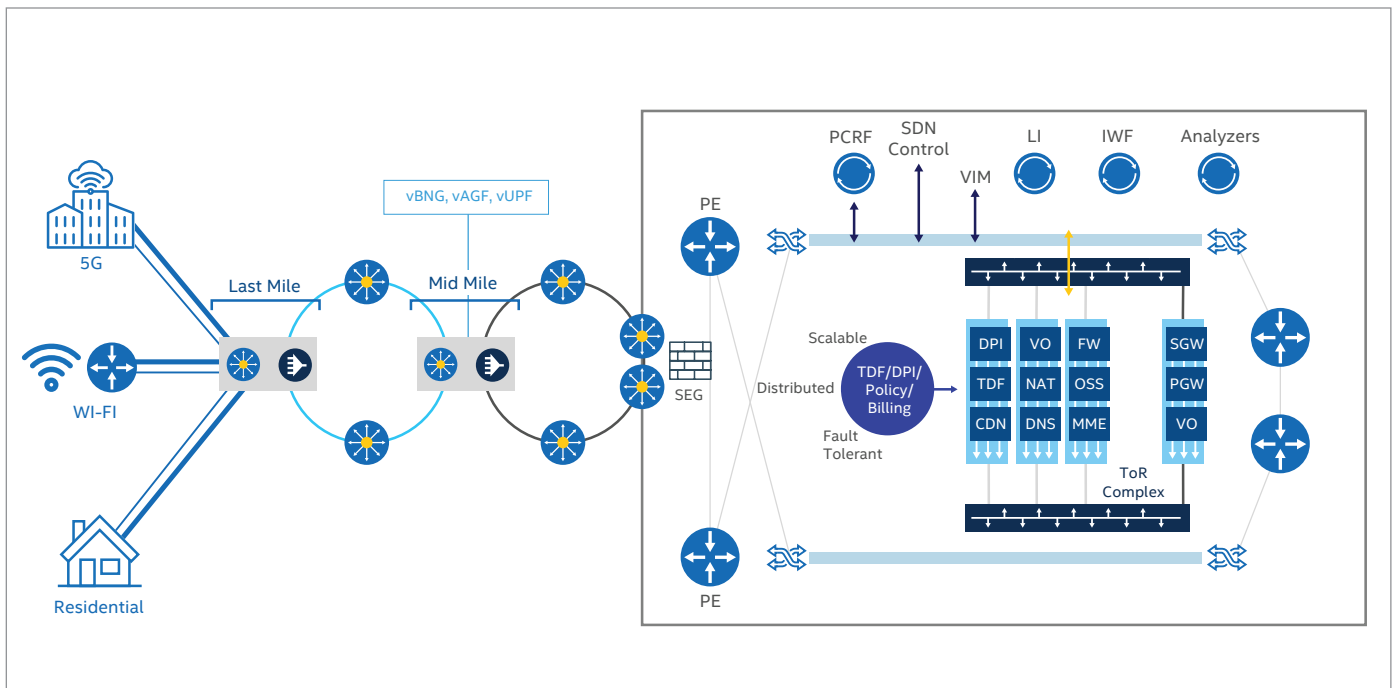


Figure 1. Next Generation Central Office

mobile applications, the mobile user plane will distribute right out to the central office location. Another emerging transformational force is the ongoing upgrade of the broadband wireline infrastructure, through which CoSPs are continuing to roll out fiber-optic cable in fiber to the home or curb (FTTH/FTTC) scenarios. At the same time, legacy copper central offices (COs) cable plants are “melting” and being replaced by fiber-fed COs primarily based on passive optical network (PON). These facilities have much longer reach than the copper COs, so that the customer base can be served by fewer distributed CO facilities. Therefore, we need fewer distributed COs to serve a given population.

This all leads to a new convergence happening at the edge of the network, whereby CoSPs across all geographies are exploring fixed-mobile converged infrastructure (FMC) at the network edge.

This (now emerged) concept of the Next Generation Central Office (NGCO) entails a fiber-rich location, where both fixed and mobile traffic can be terminated. This concept enables a more software-centric approach, which will allow operators to quickly deploy and offer new services with a more flexible architecture. The NGCO concept is intended to function as a local edge compute location toward enabling both edge NFV (for both fixed and mobile) and new Telco-hosted 5G services in a smaller form factor and power footprint than the traditional individual centralized model would enable.

Intel-based NGCO mini data center architecture is shown in Figure 1 above. This comprises industry standard servers’ storage and switching infrastructure.

This is a specific locational (central office) instantiation of edge computing and emerging paradigm, where edge breakout and computing services are performed at the “network edge”. The key network functions (e.g., BNG, UPF AGF) are virtualized or containerized and run together on general-purpose compute platforms alongside of edge

applications. NGCO enables the co-existence of VNFs and edge applications on shared general-purpose Intel® Xeon® processor-based compute platforms. This consolidation process results in converged network and service platform, which entirely transforms the approach and ecosystem used to deploy network access components. This new approach is what is known as virtualized converged access.

Cloudification of the operator’s infrastructure implies that the general-purpose nodes used to host VNFs/CNFs can now host other type of application workloads not just network related.

A key design point for edge computing is that, despite this multi-tenancy, the edge architecture must be properly orchestrated to guarantee that networking services never starve, never fail and are never impacted by other services running on the same platform. Network traffic and network function workloads needs to be deterministic.

In the following sections we will describe the requirements and methods the converged edge reference architecture needs to provide to satisfy this premise.

Edge computing allows for ultra-low latency response time and enhanced bandwidth availability in comparison to conventional centralized cloud computation models. Additional constraints may apply for compliance reasons. For example, where data is mandated (regulation) to remain within a certain location or where data must be processed in-situ because the cost of transport or transport time are prohibitive.

Applications that will be deployed at the edge include, but are not limited to, traditional network functions, connected self-driving cars, video surveillance, IoT analytics, video encoding, video analytics, speech analytics or retail services, among others.

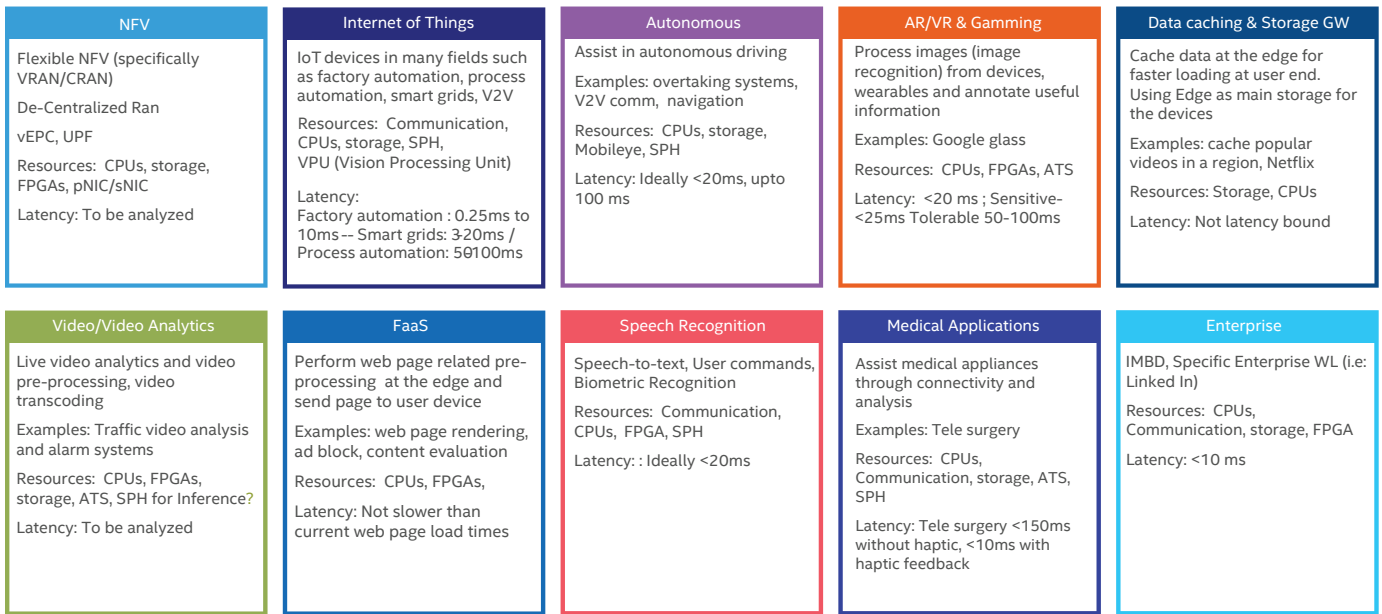


Figure 2. Use Cases for Edge Computing and their Associated Requirements

Figure 2 illustrates some of the most representative use cases that can take advantage of edge computing attributes and architecture as described above.

These services will have varying requirements in terms of:

- a. Priority (performance) and QoS – For example, traffic for an autonomous car application will have higher priority than a temperature sensor in terms of response times requirement.
- b. Reliability and resiliency – Some input streams need to be acted upon and the traffic routed with mission-critical reliability, whereas with other input streams it may be tolerable to accept an occasional failure. This instance will depend on the application.
- c. Power, cooling, and form-factor constraints.

In addition to the ability to combine both fixed and mobile traffic backhaul routing, the NGCO location therefore offers the following distinct advantages:

- a. Ability to serve and respond to multiple applications (object tracking, video surveillance, connected cars, etc.) in real-time or near real-time via mobile or fixed network access.
- b. Ability to meet ultra-low latency requirements for these applications via fiber or enhanced 5G.
- c. Ability to reduce backhaul traffic by delivering locally cached content (i.e., CDNs) and edge storage optimization.

These advantages enable a whole new class of applications (VNFs, edge applications) which cannot leverage conventional cloud computing due to latency or other requirements identified above.

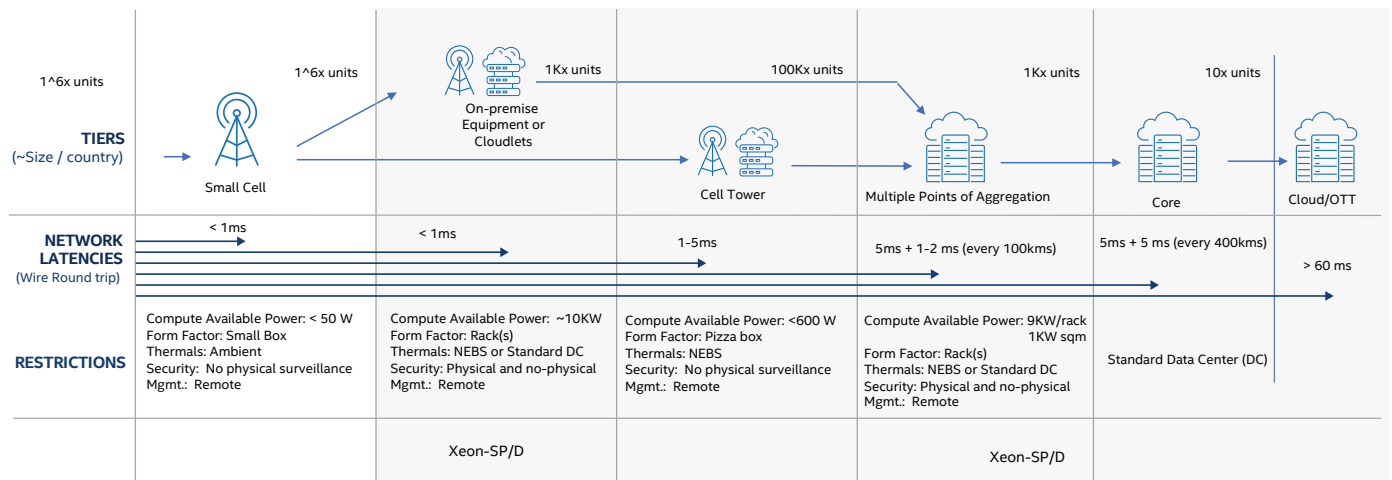


Figure 3. Latency Distribution v Platform Technologies

From an edge placement perspective there are three main different key performance indicators (KPI) that will drive the adoption of edge clouds:

- Latency is the KPI that is clearest in terms of edge definition. Physics dictates that the speed of light is ~300,000 km/s and transmission on the wire is ~2/3 of that. Thus, if some services require response latency sub 4 milliseconds, they cannot be further than ~150 km/s from the device. Thereby, for some of the workloads (i.e.: IoT) the unique edge location that will work is the base station (excludes wireline access) and central office, one step back. See Figure 3 below for the impacts of latency on platform requirements.
- Data Privacy, sovereignty and sensitivity is a compliance related value proposition that will dictate that some of the services can only reside on certain locations of the edge. For example, in the healthcare segment, some hospitals may want to host and share some services at the edge, but without having the data crossing certain boundaries of the infrastructure (i.e., equipment on premise, central office, etc.).
- Backhaul traffic data savings can be achieved by filtering traffic at the different edges of the network to reduce OPEX. In this case, filtering may happen at any of the different potential edges of the infrastructure. For example, video surveillance can be processed in the central office to identify what are the images to be sent to the cloud or processed locally.

### Objective of the stc Proof of Concept

To prove out and address the NGCO concept stc, Intel and Nearby Computing embarked on a proof of concept with interested commercial partners.

The goal of the project was to demonstrate the feasibility of deploying a completely virtualized fixed and mobile converged edge solution in the central office location utilizing the Intel converged edge concept. The ultimate target of the deployment is to utilize the 3GPP and BBF newly introduced AGF node to achieve the fixed mobile convergence at the functional level in higher granular way that share the resources between the fixed and mobile services.

In particular, the intent was to show that a general-purpose platform composed of servers powered by Intel Xeon processors can host commercial-grade VNFs and edge applications seamlessly, leveraging leading edge orchestration technology, allowing for service and network consolidation, which is a key element for TCO reduction for edge computing.

See Figure 4 below:

The use case proposed was the integration of the following key components:

**VNFs:**

- Virtualized 5G core UPF (Cisco)
- Virtualized BNG (Cisco)

**Edge Applications:**

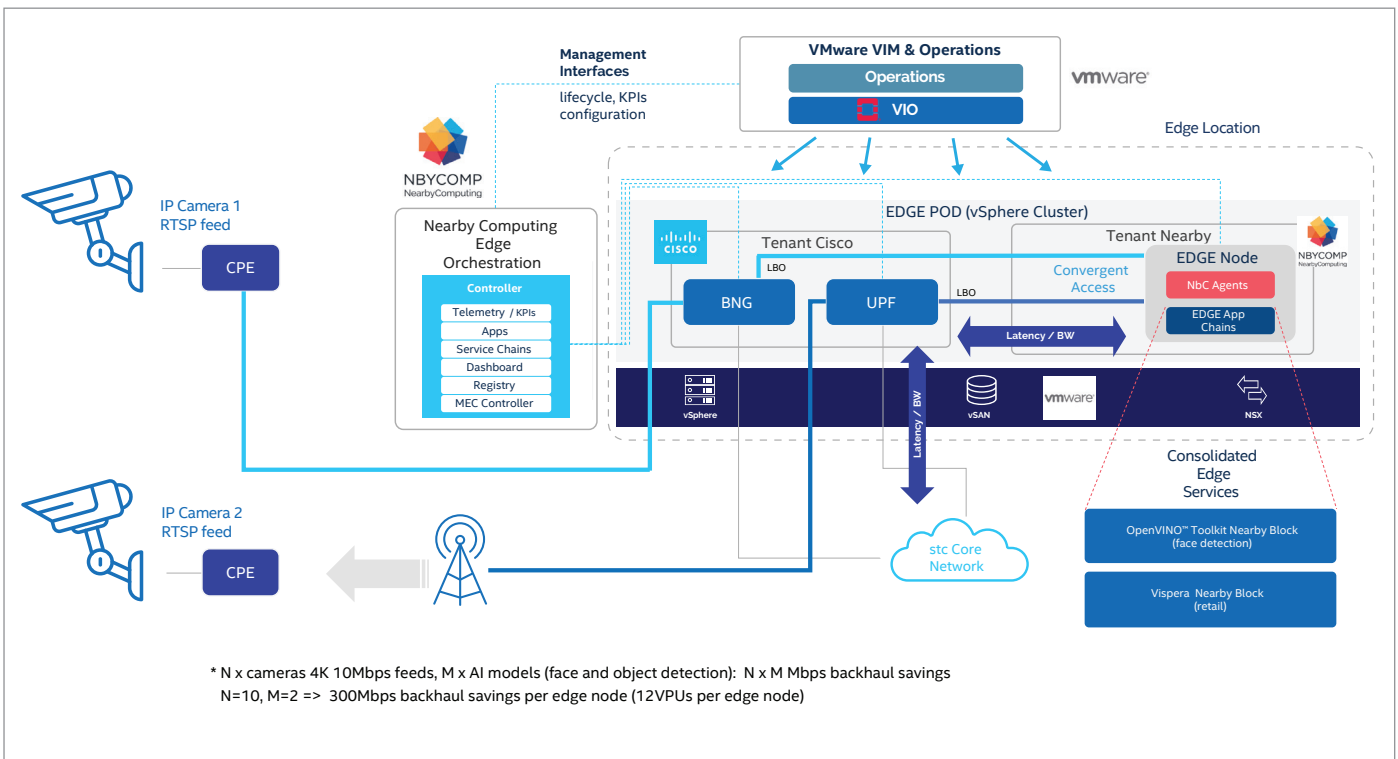


Figure 4. stc Use Case Description

- Video Analytics / Face Detection (based on Intel® Distribution of OpenVINO™ toolkit)
- Video Analytics / Retail (Vispera )

**NVFI:**

- OpenStack (VMWare VIO)

**Orchestration:**

- Nearby One solution (Nearby Computing)

**Servers:**

- Heterogeneous mix of servers powered by Intel® Xeon® Scalable processors

**Network:**

- Cisco ACI Fabric to manage the networking elements

The goal was to deploy all these components together and test their interoperability. The following Figure describes the main target components of the architecture and their main interconnection.

The following technology enablers were of relevant interest for the fulfillment of the PoC:

- The Virtualization of all software components on Intel-powered servers using a scalable NFVi
- Convergence of fixed and mobile access technologies using the Cisco VNFs for UPF and BNG all deployed on the same compute platform enabling both fixed and mobile breakout.
- Consolidation of services: VNFs and edge applications running on the same general-purpose hardware
- Orchestration and automation of all components

- Demonstration of two locally hosted edge applications where are automatically deployed and configured for the existing environment

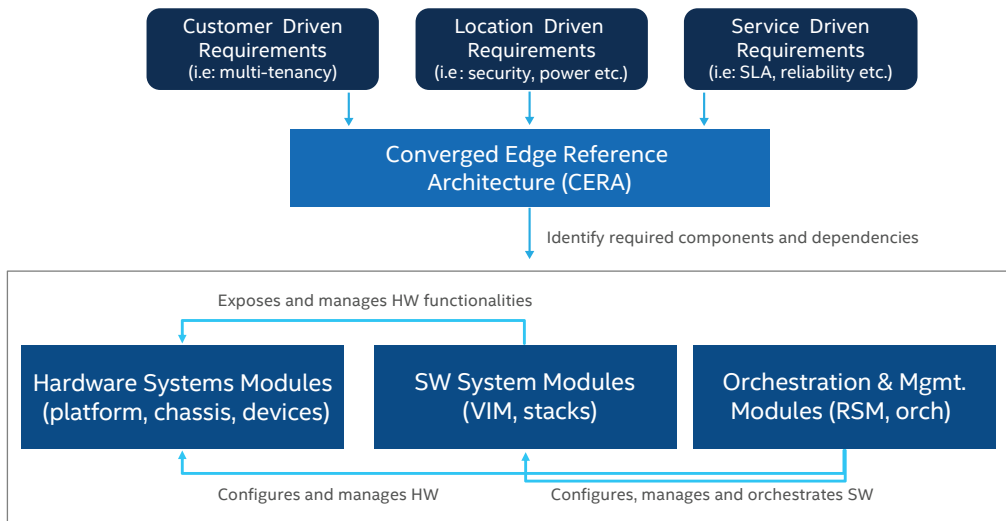
**The Outcome**

- Convergence of access: The UPF and the vBNG components were successfully deployed enabling the service local breakout of both fixed and wireless subscribers to the local services.
- Automation of all deployments: The Nearby One solution allows for VNF orchestration, and deployed both the VNFs and the services onto the compute NFVi
- Multi-Tenant Infrastructure Deployment: The Virtual Infrastructure deployment simulated for the PoC was multi-tenant , and included the following main components:
  - Core Site where three main components were run:
    - VMware VIO and vCenter components
    - 5G core control plane
    - Nearby Orchestration Platform
  - Two Edge locations running:
    - VNFs for the 5G core user plane (Cisco tenant)
    - Edge applications (Nearby Computing tenant)
    - Edge node management agents (Nearby Computing tenant)

**Intel's Approach to the Converged Edge**

	End-to-End Services										
	AI for Imaging	Content Distribution	VNFs	NLP	KVS	AR/VR					
Service Orchestration	K8s based Multi-Edge Orchestrator			K8s based Local Edge Orchestrator			K8s based Local Edge Orchestrator				
VIM	KVM	OpenStack / VMWare		ACRN	Containerization tech (VM or bare metal)						
Infrastructure	Infrastructure Telemetry					SD-WAN					
Network Infra	NIC	Smart NIC	Ethernet Switch	CXL switching	NZTP	iPXE					
Fframeworks	EdgeX	OpenVINO™ Toolkit		Edge Insights for Industrial			Edge Control Stack				
Managed Network Services	CUPS vEPC User Plane (PGW-U / SGW-U)			CUPS vEPC Control Plane (MME / HSS / SGW-C / PGW-C)				5G FWA – Fixed Wireless Access			
	Full vEPC (MME / HSS / PGW / SGW)			vBNG	FMC – AGF - FMIF	Ethernet / MPLS		vRAN			
Access	LTE / 5G NR (ORAN/FlexRan)		Fiber	WiFi / Industrial WiFi		LoRA	Ethernet	IIoT protocols			
Form Factors	Smart Camera (VPU)	No AC Cooling	AC/Liquid Cooling	Industrial Fanless	AI Appliance (VPU)	Edge Servers	Edge DC Servers	Data center Servers			
SW Platform	Security (Isecl)		Quality of Service		Acceleration	Management	Telemetry				
xPUs	GPU	AI Accelerators	FPGA(s)	Low power CPU(Atom)	Client CPU (i3, i5, i7)	VPU/HDDL	Intel® Xeon® D Processors	Intel® Xeon® E Processors	Intel® Xeon® Scalable Processors		
Platform and xPU features	Data protection (MKTME)		Workload Isolation (SGX)	Edge Attestation (TPM, TXT)		Edge Quality of Service (RDT)		AMT	TSN		

Figure 5. Intel's Converged Edge Reference Platform



**Figure 6. Steps To Configure and Deploy On Modular Edge**

## Architecture

The ecosystem of technologies enabling the edge computing paradigm is broad and covers many aspects including hardware components, service orchestration, standards and protocols, applications, virtualization technologies, and others. Figure 5, illustrates the set of technologies that Edge Computing must address for different verticals. The diversity of these technologies can pose a huge challenge for network operators, trying to select the correct combination of solutions.

The proposed edge architecture helps to alleviate this challenge and is based on three fundamental aspects combined to create a unique end-to-end ecosystem:

- It is based on standards and interoperable components to enable integrating, testing and validating different technologies from different vendors. This enables a very rich and competitive composable architecture that is not coupled to a single technology, it scales based on customer requirements
- It unifies different verticals into a single edge architecture (Telecommunications, Internet of Things, Enterprise and government) which traditionally have been implemented with different technologies and architectures. Thanks to this convergence, the overall system can be managed in a more efficient manner drastically reducing the total cost of ownership.
- It exposes a set of resources, software development kits and application programming interfaces which are easy to develop, test and deploy applications optimized and design for edge.

As a result, customers adopting the modular edge architecture can select from a number of interoperable components, hardware and software, connect them to the platform, provide resource and service requirements and let the system operate autonomously to satisfy their needs. Figure 6 illustrates the process to define the components needed to fulfill the needs of each use case evaluated by the customer:

This stc driven NGCO instantiation of converged edge architecture has a distinct value proposition that will allow stc to provide a unique solution aligned with their broader strategy that will improve:

- 1) Total cost of ownership with respect to existing solutions. Different services from different verticals are consolidated regardless of the access medium.
- 2) Provides a more flexible and malleable architecture.
  - a. Rapidly Introduce and remove new services in an access agnostic manner
  - b. Enable Zero touch provisioning and automation to a set of distributed central office locations.

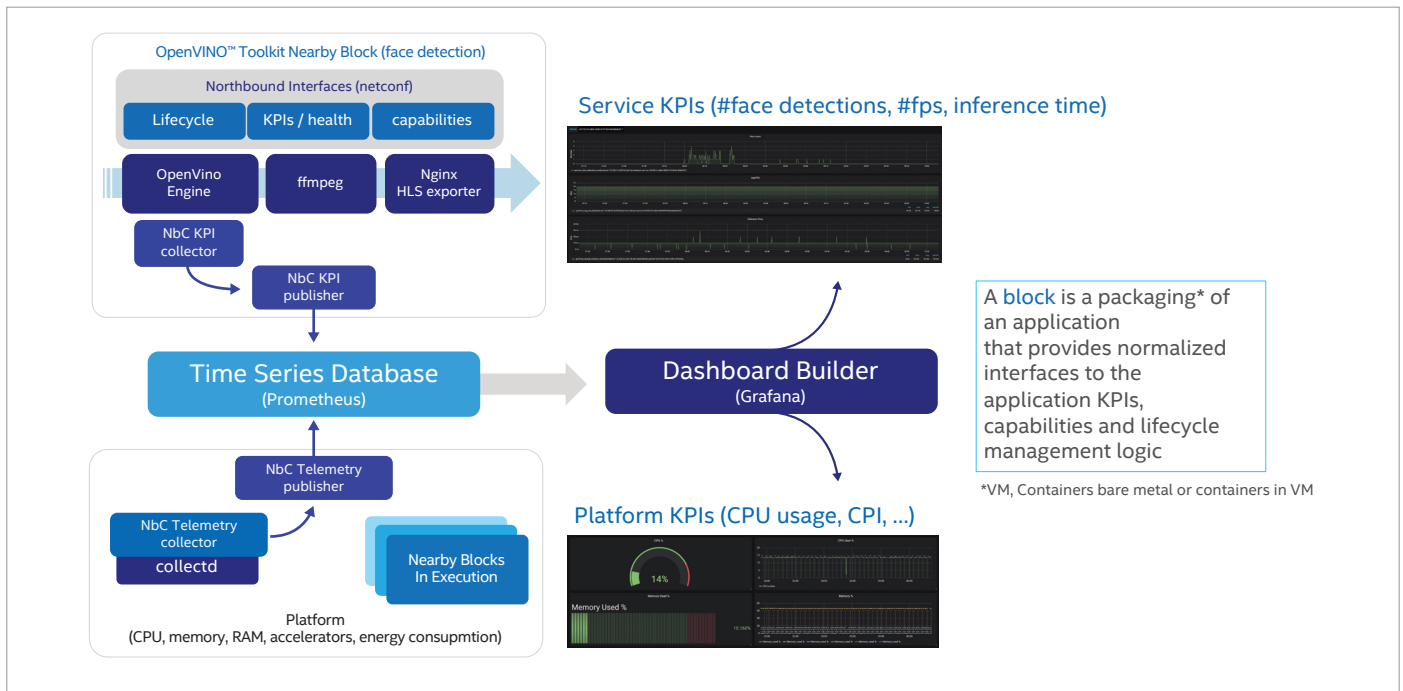
## The Nearby One Solution: Orchestration for the Modular Edge

Nearby One was used in this PoC as the edge orchestration and automation engine to manage the deployment of the services and perform end-to-end service assurance. This solution is compliant with the Modular Edge architecture previously described in this document. It is composed of:

- The Nearby Orchestration Platform is a centralized controller, providing the logic and control loops for all tasks related to the orchestration of the applications, network and compute infrastructure.
- The Nearby Blocks are distributed components that encapsulate logic and code for different application-specific functionalities, according to Intel's CERA model to extract application KPIs and expose them to the closed control loops implemented in the orchestration platform.

In this particular NGCO case, the applications were shipped as Docker containers and were deployed on emulated edge servers (VMs) that were treated as bare metal servers, while the VNFs were provided as VMs hosted on VMware in this PoC.

The video analytics application used in this PoC was powered by the Intel® Distribution of OpenVINO™ toolkit, and was onboarded and encapsulated as a Nearby Block (See Figure



**Figure 7. OpenVINO™ toolkit-based Edge application packaged as a Nearby Block**

7), which contained the application logic, as well as a set of auxiliary components that provided the means for the application to be effectively managed, including:

- Application performance KPIs for continuous SLA assessment (not only network-centric metrics)
- Application health/status KPIs for continuous lifecycle monitoring and management
- Application capabilities dynamically exposed at deployment time
- Correlation of application KPIs and platform (processor, accelerators, memory, storage) for more effective edge platform management

The solution was deployed following a full model, where all the components ran in stc central office location. The Nearby Controller was deployed in a separate location and connected by VPN provided by stc.

## Conclusion and Next Steps

This work demonstrated the feasibility of implementing a distributed modular edge architecture, based on the principles described earlier in this document allowing for the seamless convergence of virtualized services and VNFs on general-purpose Intel Xeon Scalable processor servers. Alongside an advanced orchestration engine that can coordinate the actions of VNFs and edge applications. Adopting this approach will lead to business value realization at the heart of its strategy to use the technology as an enabler to offer new digital services, leveraging its existence as a leader for the telecommunication service provider.

The PoC also proves that convergence of access technologies (Fixed/Mobile) is possible and can lead to multiple business benefits. A real infrastructure consolidation can be realized specially in the situations at which the same traffic is

swiveling between fixed and mobile nature. The pandemic situation is a clear example of how the same applications were pushed to use the fixed network instead of the mobile one due to stay at home situation.

The sudden changes of the access traffic makes central office locations the optimum ones to achieve the (fixed/mobile) infrastructure convergence. Add to this the cloud native and flexible edge approach which enables service deployment or 3rd party business models.

The three key enablers of the platform are as follows:

### Modular Intel® Xeon® Scalable Processor Architecture

- Enabling a faster time to market. Requirements are rapidly translated to modules to system architecture
- Enables an existing eco-system of services and network applications to run on the same deployment stack

### Scalable Multi-Tenant Solution stack

- Multiple services can run simultaneously and independently reducing the TCO

### Automated Manageability

- Automated deployments and operations
- Simplified services deployments
- Combine different ISV onto the same platform

The PoC utilized each of the components above which enables new business models and allows stc to utilize their existing central office locations enabled with an Intel CERA based architecture to their advantage and proceed to real commercial trials.

This aligns with the overall digital strategy of stc towards digitizing the economy and the provision of new digitizing services.

## Glossary

3GPP	Third Generation Partnership Program	LI	Lawful Intercept
AMT	Active Management Technology	LTE	Longterm Evolution
AR/VR	Augmented Reality Virtual Reality	MA	Mobile Access
ATS	Adaptive Transport Stream	MENA	Middle East And North Africa
BBF	Broadband Forum	MPLS	Multi Protocol Label Switching
CAGR	Compound Annual Growth Rate	NFV	Network Function Virtualization
CDN's	Content Distribution Network	NFVI	Network Function Virtualization Infrastructure
CERA	Converged Edge Reference Architecture	NGCO	Next Generation Central Office
CNFs	Container Network Functions	NIC	Network Interface Card
COs	Central Offices	OPEX	Operation Expenses
CoSPs	Communication Service Providers	PCRF	Policy And Charging Rules Function
CPU	Central Processing Unit	pNIC	Physical Network Interface Card
CRAN	Cloud Radio Access Network	PoC	Proof Of Concept
CSP	Cloud Service Providers	PON	Passive Optical Network
CUPS	Control And User Plane Separation	QoS	Quality Of Service
CXL	Cascade Lake Cpu	RSM	Resource Manager
DPI	Deep Packet Inspection	RTSP	Real Time Session Protocol
DSP	Digital Service Provider	SDN	Software-Defined Networking
FA	Fixed Access	SEG	Secure Email Gateway
FMC	Fixed And Mobile Convergence	SNIC	Smart Network Interface Card
FMIF	Fixed And Mobile Interworking Function	SPH	Sphere
FPGA	Field Programmable Gate Array	TCO	Total Cost Of Ownership
FTTC	Fiber To The Curb	TDF	Traffic Detection Function
FTTH	Fiber-To-The-Home	TSN	Time Sensitive Networking
HDDL	High Level Dialog Description Language	V2V	Vehicle To Vehicle
ICT	Information And Communications Technology	vAGF	Virtual Access Gateway Function
IMDB	Internet Movie Database	vBNG	Virtual Broadband Network Gateway
IoT	Internet Of Things	vEPC	Virtual Evolved Packet Core
IP	Internet Protocol	VIM	Virtual Infrastructure Manager
IPXE	Preboot Execution Environment	VNFs	Virtual Network Functions
IWF	Interworking Function	VPU	Vision Processing Unit
KPI	Key Performance Indicator	VRAN	Virtual Radio Access Network
KVM	Kernel Virtual Machine	vUPF	Virtual User Plane Function



Intel technologies may require enabled hardware, software or service activation. No product or component can be absolutely secure. Your costs and results may vary. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation, Intel, the Intel logo, Intel Xeon, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

\* Other names and brands may be claimed as the property of others.

© 2021 Intel Corporation

Printed in USA

345766-001 US

♻️ Please Recycle

0521/BB/ICMCF0/PDF