5G, Infrastructure View

Rodion Naurzalin, Lead Architect, HPE Edge Compute

Declan O'Boyle, Platform Solutions Architect, Intel Network Platforms Group

Contents

1 Use cases and architectural initiatives of 5G

With 5G standards rapidly evolving, deployments of 5G are evolving from a planning exercise into field trials and pilot deployments. A Number of industry initiatives have been investigating business requirements, use cases and architectural approaches of 5G, for example:

- NGMN Alliance published [1] their view of 5G and associated use cases, describing the following categories: Broadband access in dense areas, Broadband access everywhere, Higher user mobility, Massive Internet of Things, Extreme real-time communications, Lifeline communications, Ultra-reliable communications, Broadcast-like services.

- METIS project resulted [2] in a paper, initially describing five scenarios (“Amazingly fast”, “Great service in a crowd”, “Ubiquitous things communicating”, “Best experience follows you”, “Super real time and reliable connections”), followed by a number of test cases and industry examples. Based on insights of METIS, the METIS-II project moves forward defining technology requirements in more details

- ITU-R defined [3] main usage scenarios and role of IMT2020, classifying those into “Enhanced Mobile Broadband”, “Ultra Reliable and Low Latency Communications”, “Massive Machine Type Communications” categories

- 3GPP, starting with more than 70 use cases in TR22.891, narrowed them down to four primary use case categories of “Massive IoT”, “Critical Communications”, “Enhanced Mobile Broadband” and “Network Operations”, further defined in TR22.861-864

The uses cases above illustrate the uniqueness that 5G will bring as compared to previous generations of mobile networks – wide diversity of requirements which are sometimes contradictory to each other (e.g. massive scale with wide coverage vs. enhanced throughput, high level of subscriber mobility vs. static and nomadic use cases). To address this challenge, a new set of functional requirements and architectural approaches to 5G has been generally agreed on by the industry as shown below[4],[5],[6], [7], [8]:

- **Separation of control and user plane**
  Practical experience of running and evolving a mobile network shows that scaling properties of control plane and user plane can be very different even in a network dedicated to a fairly homogeneous use case (e.g. Mobile Broadband). In a 5G world, serving traffic with very different and dynamically changing signaling vs. payload profiles is the norm. Moreover, certain 5G use cases impose fundamentally new levels of requirements on efficiency of control plane and user plane which are independent of scaling. High reliability, low latency communication and enhanced mobile broadband use cases define a need for not only functional separation and independent scaling, but optimal and dynamic placing of signaling plane and user plane functions in a network, based on operator policies and demands of a concrete use case

- **Network slicing**
  One of the most anticipated functionalities of the 5G network is the ability to programmatically create, modify and delete network slices. This allows an operator to provide different sets of services, QoS/QoE profiles, or serving different subscriber groups using the same physical 5G infrastructure. A network slice provides end-to-end network service including access, core functions, associated policies, etc – all optimally placed and configured to satisfy the requirements associated with that slice.

- **Split architectures**
  The overall theme of flexibility and optimized placement of network functions in 5G defines new architectural approaches to next generation Radio Access Networks. What previously was seen as a monolithically implemented and deployed access domain is now evolving capabilities to be

split along well-defined interfaces (e.g. the CU-DU separation, RAN control-user plane separation), as well as to place core network functions and end user applications adjacent to radio access domain at the edge of network

- **Access convergence**
  A 5G network is expected to support multiple access types (e.g. Fixed Broadband, WLAN) based on operator policies and needs of a concrete use case. This results in high diversity and density of network functions placed throughout the 5G network

- **Virtualization**
  Today Network Functions Virtualization became a well-proven technique in securing flexibility, reliability, scalability and cost efficiency for core elements of a mobile network. With new split architectures in RAN as well as flexible placement of 5G core elements from edge to core, virtualization is likely to become a key enabler in building a highly efficient and flexible 5G network edge to core

The new 5G use cases and architectural initiatives to support them, as well as a trend for overall function virtualization have a common theme – a need for universal IT-grade, software-defined infrastructure, spanning well beyond telco datacenters to the edges of 5G network.

2 Preparing infrastructure for 5G

A mobile network is a unique example of massively distributed infrastructure: a mobile operator's radio access network can span tens of thousands of physical sites. These are supported by thousands of traffic aggregation sites and other points of presence (e.g. Regional POPs, Central Offices, etc.). However in traditional mobile deployments, the network functions at the access sites and other PoPs are static and in many cases implemented as special purpose compute appliances.

The decoupling of a network function from underlying hardware, automating lifecycle management of the network functions as well as orchestrating services using Network Functions Virtualization (NFV) technologies has, to date, been limited to network functions hosted in a datacenter environment. The 5G architectural developments as well as a maturing NFV ecosystem can extend the concept of flexible dynamic placing of network functions far beyond datacenters – to what has traditionally been considered "network edge" (e.g. Central offices, branch offices, traffic aggregation points) and even beyond that – to the access network (RAN sites, backhaul aggregation points).

As we noted above, many 5G use cases require flexible steering and termination of user plane and control plane to enable low latency and high reliability. These principles of flexible traffic steering and early break-out from mobile network are addressed by Multi-Access Edge Compute (MEC), enabling the placement of application workloads throughout network fabrics, edge to core, as it fits requirements of a particular use case. Examples of such highly-distributed applications embedded into mobile network might be video recognition and computer vision, AR graphics rendering and AR streaming, control of distributed city infrastructure or coordination of autonomous cars.

It is worth noting that MEC is not unique to 5G – it is just as viable for LTE and can be implemented today with multiple architectural approaches (e.g. early traffic breakout by placing components of LTE packet core at the edge, or introduction of a function tapping into the S1 interface between eNodeB and packet core).

MEC is not the only point of convergence between 5G and previous generations of mobile networks – communication service providers need to continue serving their existing subscriber base while introducing new innovative services. This means that 2G, 3G, 4G and 5G need to continue to co-exist. That is confirmed by 3GPP studies of 5G deployment options – introducing up to 7 options of 5G and 4G coexistence, including those of 4G/5G dual connectivity. Many of the 5G architectural approaches (such as control-user plane separation in core and RAN, virtualization and dissagregation of RAN) can be adopted for implementations of legacy technology as well.

Next generation flexible x86 infrastructure serves as an important anchoring point for the evolution of existing network technologies, introducing 5G technology and addressing new business cases with MEC. Figure 1 shows a deployment option, where evolved architecture of LTE, 5G components and applications (enabled by MEC) coexist on the same infrastructure, securing the highest return on investment into new technologies, as well as delivering flexibility to further evolve the architecture to service new, yet unknown, use cases.

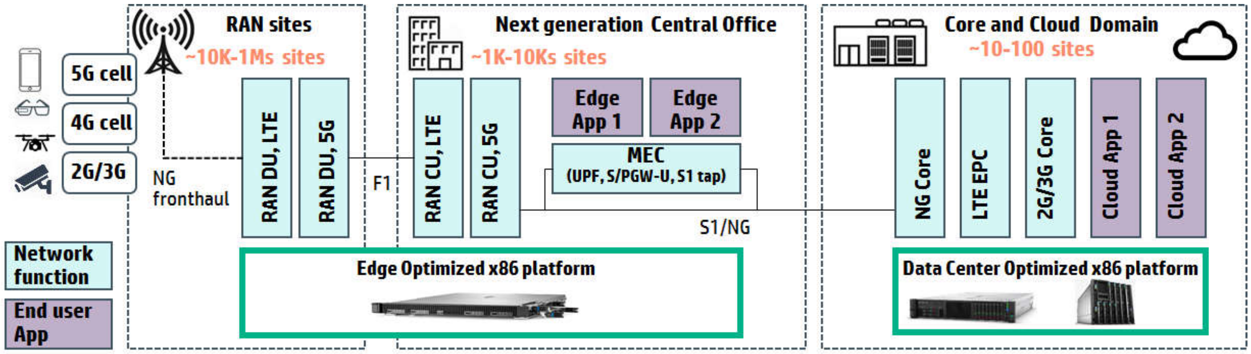

Figure 1 Example of 5G deployment leveraging a converged x86 infrastructure

Figure 1 illustrates an example of how VNF and Application workloads might be placed within common 5G infrastructure. Flexible dynamic placing of VNFs as well as placing non-telco Applications from edge to core of a 5G network require fundamentally new capabilities of underlying infrastructure.

One of the most disruptive aspects of the 5G deployment example illustrated on Figure 1 is a need for massively distributed generic compute resources at the edge of a network. With diverse and dynamic workloads placement, the edge infrastructure should be abstracted and automated in a similar fashion as currently centralized cloud environments are. With that in mind the 5G edge infrastructure is sometimes referred to as “Massively distributed edge cloud” or simply “Edge Cloud”.

While edge clouds have to provide similar abstraction and automation as centralized cloud environments, they differ significantly in how they are deployed and operated. Edge Cloud infrastructure is massively distributed and operates in physical environments different from what is available in datacenters. Furthermore, it is expected to host a growing number of diverse workloads (VNFs and applications), some of which require specialized services from underlying infrastructure (such as hardware acceleration techniques).

These distinct characteristics of edge clouds place unique requirements on the underlying general purpose compute infrastructure. In the following sections we describe how x86 general purpose processing addresses the needs of the 5G network and present the main functional requirements towards an edge-optimized general purpose compute platform.

2.1 Role of General Purpose Processing in 5G network

Proliferation of NFV at telco datacenters made General Purpose Processing one of the key enablers of an efficient, reliable and flexible telco network.

In the following sub-sections we provide more specifics of how GPP is used to address the architecture requirements and features of 5G which we discussed above.

2.1.1 Flexible deployment of vRAN

Traditionally, the base station was a fairly static building block of RAN – with monolithic appliance-based baseband modules which processed a full RAN protocol stack, and connected to radio modules generating radio signals.

However the continuing introduction of new spectrum bands, requirements for simultaneous support of heterogeneous combinations of these bands on different sites, advanced antenna techniques as well as overall densification of RAN networks make the economics of scaling such traditional base stations unattractive. Massive-MIMO deployments further illustrate the necessity to scale the baseband processing envelope. Moreover, centralization of baseband functions among the large number of base stations promises enhanced abilities in interference control, advanced resource scheduling and mobility while pooling resources can improve the overall CAPEX/OPEX. However, addressing the scaling challenge and attaining the benefits of functional centralization and resource pooling comes with a trade-off.

Centralizing baseband functions for the lower layer protocols of a baseband requires a low-latency and higher bandwidth (relative to S1 interface) on the front-haul between the radio sites and centralized location. This essentially implies that optical fiber is necessary (although research activities are ongoing into the viability of single-hop microwave links).

Distributing the baseband processing remains an attractive option in scenarios where the transport infrastructure cannot meet the latency or bandwidth requirements of a centralized approach. The very low-latency sensitive use-cases will require base-band processing near the radio.

As a result the next generation deployments will consist of a mix of centralized and distributed baseband processing. Distributed BBUs will be deployed where the transport infrastructure does not satisfy quality and bandwidth requirements necessary to transport the front-haul traffic. Additionally base-band processing closer to the end-user may be essential for very low latency applications.

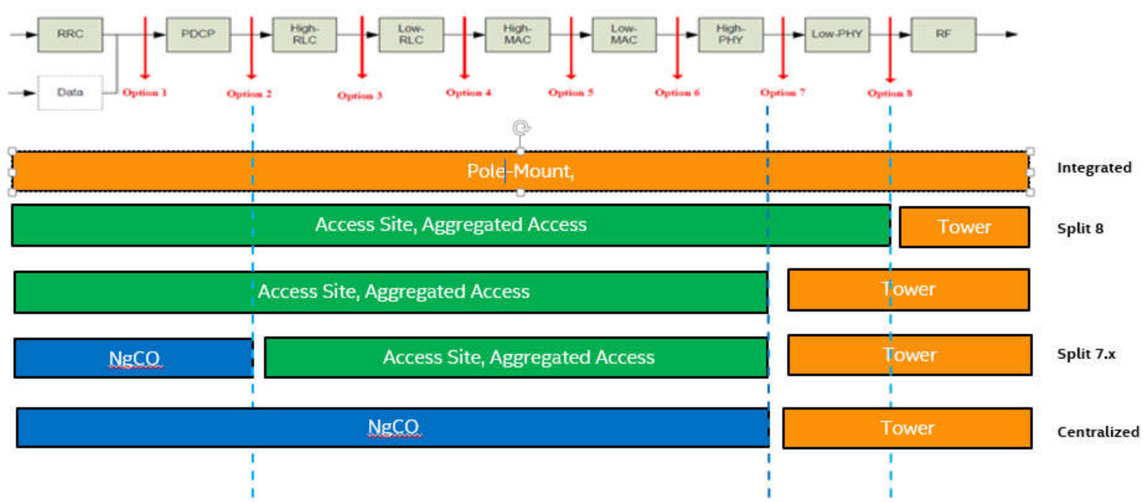

Figure 2 BBU split processing options at various network locations

In this context the value of virtualized RAN functions are more apparent. As Figure 2 illustrates, the RAN SW stack can be implemented as a set of vRAN modules that can be flexibly deployed at various network locations on x86 platforms – for both high level and low level protocol layers. Consequently, a consensus is emerging that all layers down to the upper-L1 can be effectively implemented using virtualized functions that leverage the benefits of NFV.

Today, the upper layers of the RAN stack (PDCP and RLC) are already being deployed at centralized locations and execute as a VNF running on standard x86 infrastructure. Centralizing and deploying the lower layers of the baseband stack (e.g. L1/L2 or portion of those) on x86 compute is the logical next step – provided that challenges associated with real-time requirements can be addressed.

Putting lower layer stack on x86 architecture requires a number of critical enabling technologies: software enhancements to standard Linux distributions to provide real-time kernel functionality, low latency processing and efficient virtualization implementations. Some examples are:

- Technologies such as DPDK and Dynamic Device personalization (DDP) enable competitive packet processing on x86 architecture and route the data to the CPU cores in efficient manner.

- Dynamic Device Personalization on Intel® Ethernet 700 series controllers enables dynamic reconfiguration of the packet processing pipeline and provides the ability to classify new packet types (as is likely to be defined by Front-haul Ethernet protocols) and distribute these packets resulting in increased system performance due to reduced packet latency and offloading packet classification and load balancing functions from the CPU Cores. This capability accelerates packet processing for emerging front-haul Ethernet protocols such as IEEE1914 RoE and eCPRI.

- The introduction of Intel® Advanced Vector Extensions 512 (Intel® AVX-512) on the Intel® Xeon® Scalable processors and new Intel® Xeon® D processors improves the data throughput compared to previous generation and results in significant performance gains for the physical layer signal processing functions of the 4G & 5G RAN stack.

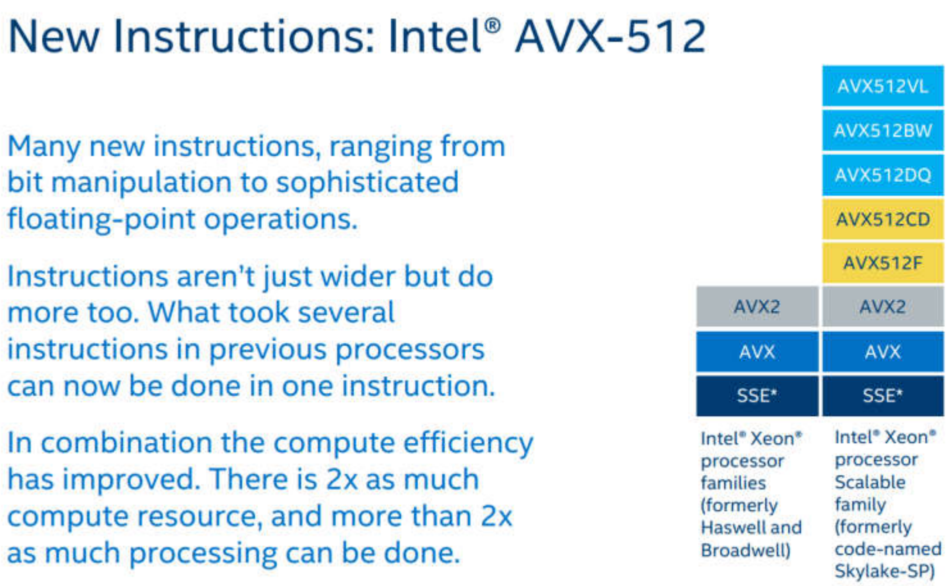

Figure 3 Building efficient 5G-NR base stations with Intel Xeon Scalable Processors

The common thread is that the platform architecture remains consistent, provides efficient common virtualization implementations across the various platforms that supports common software modules.

2.2.2 Independent HW & SW Upgrade

Another key promise of General Purpose Compute in 5G network is higher operational efficiency. Purpose built hardware appliances for baseband processing are traditionally deployed for an average of 5 years, hence during the network design and development phase, headroom is usually built in to accommodate future SW upgrades. The software stack is specifically designed and optimized to execute on the specific hardware and tightly coupled with it. In previous communication technology generations (2G, 3G, 4G) it has been relatively easy to predict future enhancements features that the appliance hardware may need to accommodate. However given the expected diversity of 5G applications, this has become impossible.

Software compatibility across previous and next generation Intel Architecture provides a decoupling of the hardware from the software. Next generation hardware can be deployed independent of major software upgrades and transforms the risk and cost traditionally associated with hardware upgrades as well as eliminates any vendor lock-in associated with installed base of proprietary hardware appliances.

2.2 What is a 5G edge-optimized platform

Having discussed what a 5G Edge Cloud needs to be able to do, let us now discuss what it is. Specifically, 5G Edge Cloud needs to have the following features.

- **Compute Density**

With a growing number of VNF and Application workloads co-hosted at the edge of the 5G network, the underlying infrastructure should provide maximum compute horsepower given the limited footprint of an edge site (especially important at deep network edge such as RAN site)

- **Universal compute edge-to-core**

5G Edge Cloud infrastructure hosts workloads very different in their nature, co-hosting mobile network VNFs as well as end user applications on the same infrastructure. With flexible and dynamic placement of network functions and end-user application in 5G, the Edge Cloud should enable migration of the workloads between different edge clouds and from edge to core and back as needed. That requires the underlying general purpose edge compute platform to provide compute resources consistent with those available in the cloud (e.g. standard x86 CPU resources, GPU resources). However unlike the cloud where certain types of resources (e.g. GPU compute resources, AI accelerators) are pooled in their own appliances, edge compute platforms should provide a high level of convergence among those within the same box, while exposing them to applications in a way consistent with the cloud

- **Optimized scaling at the edge**

One of the expected attributes of 5G Edge Cloud is uneven distribution of compute resources throughout the edge domain. Dynamic workload placement requirements as well as difference in physical characteristics of edge sites dictates a need for highly modular and scalable platforms for Edge Cloud

- **Power Efficiency**

Deployment of significant generic compute capabilities throughout the edge domain of a 5G network requires careful planning of supporting infrastructure, particularly for power supply systems.

In addition to general considerations for energy efficiency, the physical environment of Edge Cloud sites is different from the typical datacenter. RAN sites and traffic aggregation hubs are usually equipped with single unreliable sources of power, supported by different forms of battery back-ups or diesel generators for cases of main failure. Deployment of additional compute platforms on these sites might lead to a re-design of power supply units, battery backup systems and generator equipment – the costs of which can easily negate any benefits of edge deployment. Thus power efficiency becomes a critical factor for edge-optimized compute platforms to be economical for 5G deployment as it determines not only OPEX of electricity consumption, but CAPEX associated with potential upgrades of power systems throughout thousands of sites.

- **Security considerations**

5G Edge Cloud enables deployment of critical functions such as components of mobile core and application logic at sites, which can’t boast the same level of physical and cyber security of a datacenter. If not properly addressed, this significantly increases the susceptibility of attack on a 5G network, therefore security of an edge-optimized compute platform become a mandatory requirement.

2.3 Operational capabilities of a 5G edge-optimized platform

Having briefly discussed how to build a 5G Edge Cloud, let us now recall that it has to operate over a large number of widely distributed physical locations, traditionally used for hosting telco purpose-built appliances (e.g. radio base station equipment at RAN sites, access routers at traffic hubs, etc). These hosting environments as well as operational practices of a telco operator for equipment placed at these sites are very different from a typical datacenter. Figure 4 shows an example of a potential 5G deep edge hosting environment, based on a typical RAN site layout.

For smooth and economical deployment and operations of general purpose IT hardware at deep edge sites, edge-optimized 5G platforms should address the environmental specifics:

- **Consumer-grade climate control**

  Climate control at a deep edge site is usually performed with consumer-grade HVAC systems which bring significant temperature and humidity deviations compared to a tightly controlled climate of a datacenter.

Figure 4 Deep edge hosting environment, based on a typical RAN site layout

- **Limited physical space**

  Deep edge sites usually provide very limited physical space for equipment. Edge-optimized general purpose hardware should be prepared for various mounting and racking options (including wall mounting and shallow racks), providing maximum compute resources per occupied footprint.

- **Limited site accessibility**

  Deep edge hosting locations such as RAN sites usually leverage part of commercial or private property, leased by a telco operator for equipment placing. In many cases that means complicated physical access procedures for maintenance purposes (e.g. site visits are subject to additional agreements with a landlord, one-off permits, limited number of entrances per period of time, or even no physical access majority of time in case of special nature of leased real estate). Limited physical site accessibility puts strict requirements on reliability of equipment placed at deep edge sites.

- **Deployment and operational practices**

  With decades of experience in building and operating massively distributed infrastructure, mobile telco operators built efficient operational processes and organizational frameworks for storing, delivering, installing and operating traditional telco appliances at the deep edge sites. Cost-efficient and smooth introduction of general purpose hardware at deep edge sites demands leveraging operator's know how and existing operational processes.
  From that perspective, edge-optimized hardware platforms should have non-operational characteristics close to traditional telco appliances.
  One example is logistical operations, optimized by a telco operator for physical delivery of infrastructure to remote sites. To be safely stored and transported while leveraging existing operational processes built around characteristics of telco appliances, the edge optimized platform should tolerate extended non-operational temperatures and humidity typical for warehouses as well as high tolerance for shocks, inevitable during transportation to massively distributed physical locations.

- **Platform manageability**

  With potentially tens of thousands of deep edge sites deployed, the underlying edge optimized x86 platform should provide means for platform provisioning, assurance and management at scale. On a platform level, that requires open and standardized interfaces for platform management, compliant with industry standards (e.g. RedFish API) and consistent with those of the data center platforms.

3 Conclusion

While pursuing major opportunities that 5G brings, telco operators need to address a number of challenges associated with cost-efficient network deployment that scale. A key enabler of efficient 5G deployment is underlying universal compute infrastructure spanning from core to the very edge of a network.

This infrastructure provides the following benefits for a 5G network deployment:

- Enables flexibility in deployment of new architectures with dynamic placement of network functions and end user applications edge-to-core as needed for particular use case
- Unlocks full potential of virtualization in new network domains such as RAN
- Improves long term efficiency of investments into 5G and overall TCO by decoupling the life cycle of network functions from underlying infrastructure

To fulfill these promises the 5G compute infrastructure should have the following properties:

- Universal general purpose compute resources that are available from edge to core and exposed in consistent way throughout whole 5G network
- At the edge of 5G networks, general purpose compute platforms should be optimized for new unique requirements to operate outside of datacenters (tough operational conditions, limited footprint, non-secure environments)
- General purpose compute platforms should support consistent mechanisms of deploying and managing the infrastructure both at edge locations and in datacenters

With the goal of unlocking the full 5G potential, Hewlett-Packard Enterprise and Intel are collaborating in building general purpose compute platforms optimized for unique challenges of a next generation mobile network from the core to edges.

Appendix A. HPE 5G edge-optimized solutions

As a concrete example of an edge-optimized solution, we discuss HPE's Edgeline platform. Hewlett-Packard Enterprise designed EdgeLine platform for unique challenges of 5G Edge Cloud. The EdgeLine platform has a unique combination of functional and operational capabilities making it a hardware platform of choice for deep edge deployments of 5G Edge Cloud.

Figure 5 EdgeLine EL1000, EL4000

- **Compute Density**

Designed as a System-on-a-Chip (SoC), EdgeLine platform reaches density of up to 64 physical compute cores per single unit chassis (with the number growing in future generations of platform), making it 2-3 times more dense comparing to a typical data-center optimized hardware. That unparalleled density makes it possible to co-locate compute intensive workloads at the limited footprint of the deepest edge (for example, co-hosting CRAN, UPF and Video Analytics workloads at the RAN site).

- **Compute tailored per application needs**

EdgeLine platform design enables diverse options of specialized compute without any loss in generic x86 compute density. It provides highest density of x86 compute, built-in GPU accelerators supporting heavy graphic workloads, dedicated accelerators such as GPU, FPGA, QAT modules supporting acceleration of network functions and end user applications. EdgeLine platform can be equipped with industry state-of-art platform for hardware acceleration of deep learning models at the 5G edge.

- **Optimized scaling at the edge**

EdgeLine platform is designed with modularity requirements in mind. The system scales from single compute node to a multi-node within single chassis by adding universal SoC-based server modules. Figure 4 illustrates how an edge platform can scale from single node to multi-node fitting requirements of particular site of 5G Edge Cloud.

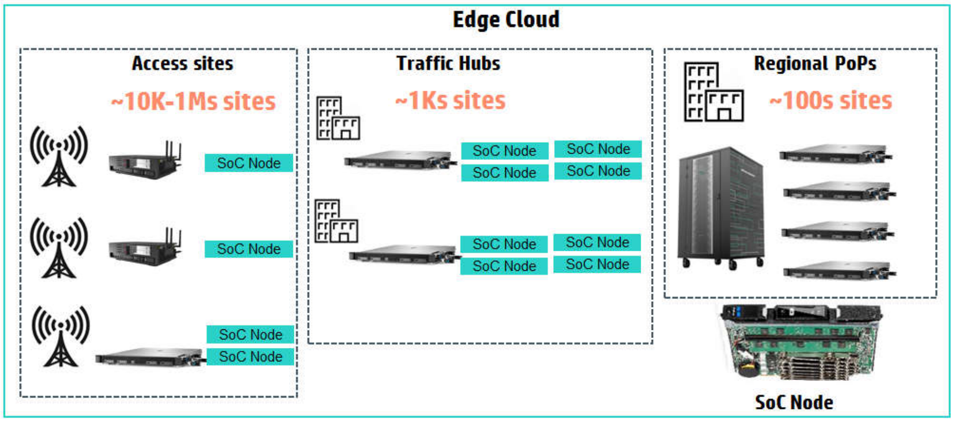

Figure 6 Edge node scaling within 5G Edge Cloud

Deepest edge sites can be equipped with single-node EL1000 chassis, co-deployed with special-purpose built telecommunication equipment present at the site (e.g. inside base station enclosure, racked with access network routing equipment, etc).

As the need for generic compute at the edge grows, multi-node edge sites located far beyond datacenters can be built as a single chassis of EL4000 enclosing up to 4 SoC Nodes, inter-node switching, HW acceleration cards in a single one-unit box, limiting physical footprint and simplifying system maintenance at the deep edge site.

In case an operator needs even more processing power at the edge (e.g. higher-capacity systems deployed higher up in edge hierarchy - in regional PoPs or branch offices of an operator), multi-node edge site can be deployed as a stand-alone micro datacenter enclosure, equipped with number of EL4000 chassis providing highest density compute at the edge.

- **Power efficiency**

Power efficiency considerations drive EdgeLine platform design from ground up. With tight components packing based on SoC architecture and power-efficient CPUs, EdgeLine platform delivers about x2 times reduction of Watts-per-Core comparing to traditional datacenter-optimized platforms. Power savings techniques of Intel Architecture ensure power-consumption is based on the actual workload at that point in time rather than a fixed power-consumption for peak use-case as typical with fixed function HW.

- **Hardware-rooted security**

EdgeLine platform implements hardware-based root of trust based on Trusted Platform Module compliant with TPM 2.0 specification. Leveraging TPM 2.0 capabilities, security can be increased by storing the encryption startup key in hardware of the server, as well as storing any other passwords, certificates, and encryption keys that can authenticate server hardware and software.

- **Operational characteristics**

EdgeLine is designed from ground up to be able to operate in harsh environments. Extended operational temperature range of EdgeLine platform is consistent with that of a typical telco appliances running on deep edge sites today. That enables EdgeLine to run in the same environmental conditions, without significant upgrades to climate control systems (e.g. installation of tight climate control enclosures) or loss of resiliency and performance.

EdgeLine platform designed to provide maximum compute density per unit as well as various mounting options. Having system service access from sides of a unit and configurable air flow, EdgeLine chassis can be mounted in various positions – in a rack, on a shelf, on a wall, making most of the space available at the deep edge physical site as illustrated on Figure 7.

Figure 7 Flexible mounting options for EdgeLine EL1000 and EL4000

To be consistent in equipment reliability, EdgeLine system is optimized to provide higher MTBF by means of tight coupling of elements within SoC and extensive end-to-end system testing.

To satisfy mass deployment requirements, EdgeLine exposes RedFish-compliant iLO interface. That allows to provision, manage and operate the platform in consistent way with data center platforms, leveraging some of the tooling and best practices.

Unparalleled density, enhanced environmental characteristics, edge optimized form-factor as well as IT-standard manageability makes EdgeLine an ideal candidate for 5G infrastructure preparation.

Acronyms

Acronym	Value
3GPP	3rd Generation Partnership Project
API	Application Program Interface
AR	Augmented Reality
BBU	Baseband Unit
CAPEX	Capital Expenditure
CPU	Central Processing Unit
CU	Centralized Unit
DDP	Dynamic Device Personalization
DPDK	Data Plane Development Kit
DU	Distributed Unit
eCPRI	Evolved Common Public Radio Interface
FPGA	Field Programmable Gate Array
GPP	General Purpose Processing
GPU	Graphics Processing Unit
HVAC	Heating, Ventilation, and Air Conditioning
IEEE	Institute of Electrical and Electronics Engineers
ITU-R	International Telecommunication Union, Radiocommunication center
LTE	Long Term Evolution
MEC	Multi-access Edge Computing
MIMO	Multiple Input Multiple Output
NFV	Network Functions Virtualization
NgCO	Next generation Central Office
NGMN	Next Generation Mobile Network
OPEX	Operating Expenses
PDCP	Packet Data Convergence Protocol
PoP	Point of Presence
Intel®QAT	Intel® QuickAssist Technology
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
RLC	Radio Link Control
RoE	Radio over Ethernet
SoC	System on a Chip
TCO	Total Cost of Ownership
TPM	Trusted Platform Module
UPF	User Plane Function
VNF	Virtual Network Function
vRAN	Virtualized Radio Access Network
WLAN	Wireless Local Area Network

References

1. NGMN, "5G White Paper", 2015, https://www.ngmn.org/fileadmin/ngmn/content/downloads/Technical/2015/NGMN_5G_White_Paper_V1_0.pdf
2. METIS Deliverable 1.5, "Updated scenarios, requirements and KPIs for 5G mobile and wireless system with recommendations for future investigations", 2015, https://www.metis2020.com/wp-content/uploads/deliverables/METIS_D1.5_v1.pdf
3. ITU-R, M2083 IMT Vision, "Framework and overall objectives of the future development of IMT for 2020 and beyond", 2015
4. 3GPP TS 22.261, "Service requirements for the 5G system", 2017
5. 3GPP TR 38.913, "Study on Scenarios and Requirements for Next Generation Access Technologies", 2017
6. 3GPP TR 23.214 "Architecture enhancements for control and user plane separation of EPC nodes" , 2016
7. 3GPP TS 23.501 "System Architecture for the 5G System", 2017
8. 5G-PPP, "5G Vision", 2015
9. https://www.dmtf.org/standards/redfish
10. ETSI GS NFV 002, Network Function Virtualization, Architectural Framework
11. https://wiki.openstack.org/wiki/Fog_Edge_Massively_Distributed_Clouds
12. ETSI GS NFV 001, Use cases
13. Reference for Intel DDP, https://software.intel.com/en-us/articles/dynamic-device-personalization-for-intel-ethernet-700-series