WHITE PAPER

Communications Service Providers
Network Edge

Creating the Next Generation Central Office with Intel® Architecture CPUs

Authors

Rory Browne,
Paul Mannion,
Eoin Walsh
Intel Corporation

Table of Contents

Abstract

This white paper describes how and why communications service providers (CoSPs) are transforming their edge networks to Intel® infrastructure, allowing them to deliver increased service velocity with improved flexibility and efficiency.

Executive Summary

In their search for greater flexibility, new business models, and lower costs, communications service providers (CoSPs) continue to adopt software defined networking (SDN) and network functions virtualization (NFV). SDN and NFV, in turn, are driving a significant change in the CoSP's approach to edge evolution of the traditional central office (CO).

SDN/NFV-centric network transformation began in in November 2012 with the operator-led advent of the NFV Industry Specification Group (ISG) at the European Telecommunications Standards Institute (ETSI).[1] Many of those initially defined NFV use cases referenced existing functions and services that could be hosted as a virtual network function (VNF) in the operator's edge network, such as:

- Virtualization of content delivery networks (vCDN)
- Virtualization of the customer premises equipment (vCPE)
- Virtualization of the fixed access network
- Virtualization of the mobile core network
- Virtualization of radio access networks (vRAN)
- Virtualization of traffic analysis, optimization, and security network functions.

Each of the above offer scenarios where there is a locational (edge) advantage to hosting the relevant VNF service in the CoSP's edge. The advantages of putting the function or service closer to the end user include improved network resilience, lower latencies and jitter, lower load on the CoSP's core networks, and the possibility for new services, business partnerships, and revenue models.

Various adjacent initiatives were started to specify and standardize enabler technologies for edge distribution, including Control and User Plane separation (CUPS) and Central Office Re-architected as a Datacenter* (CORD*). These efforts are happening for fixed networks in the Broadband Forum* and for mobile networks as part of the 5G standardization efforts in 3GPP. The combination of NFV/SDN, CUPS, and CORD greatly facilitates fixed-mobile convergence, fulfilling the promise of a converged single network approach. These initiatives will be important to 5G networks because cost pressures and the technical demands of 5G make such convergence all but inevitable.

To meet the demands of the new software-centric networking paradigm, traditional equipment can be upgraded to become more configurable and programmable. New and legacy hardware supporting dedicated networking chip programming languages like P4 (programming protocol-independent packet processors) is also becoming predominant, enabling such a programmatic capability. At the same time, the capabilities of industry standard servers are steadily improving to match the networking capabilities of hardware appliances, in addition promising much needed flexibility, code portability, and lower costs through more flexible software licensing models.

In light of this technology evolution, this paper discusses the role Intel plays in the next generation central office (NGCO). The NGCO is an SDN and NFV-enabled CoSP location that is architected for delivering agile wireline or mobile network infrastructure and associated service delivery. The VNF performance the industry is witnessing on the latest Intel® Xeon® Scalable processors reconfirms this as the optimal choice for 5G and fixed edge network evolution.

1 Telecom Network Transformation

SDN, NFV, the adoption of cloud services, and the imminent arrival of early 5G architectures are driving fundamental change in the way CoSPs plan, deploy, and manage infrastructure. One result of this network evolution will be an open platform-based ecosystem approach to the network that enables new business models and possible partnerships with over-the-top (OTT) service providers and other verticals in ways that were previously impossible.

1.1 Costs and Competition Are an Impetus for Change

Unlike in years past, today's broadband-dominated mobile traffic is becoming increasingly decoupled from the growth in revenue. Subscribers pay less per bit for service, resulting in flat revenue growth and much faster growth in the costs required to support traffic growth, such as capacity, coverage, and quality of service upgrades, outpace the projected revenue growth. Traditional communications equipment costs drop at a reported rate of 10-15 percent annually while traffic growth continues unabated at 30 percent annually.[2] Hence, CoSPs must now explore alternative approaches to reduce costs and build out new services, allowing for the possibility to quickly and efficiently trial/grow new service and revenue streams.

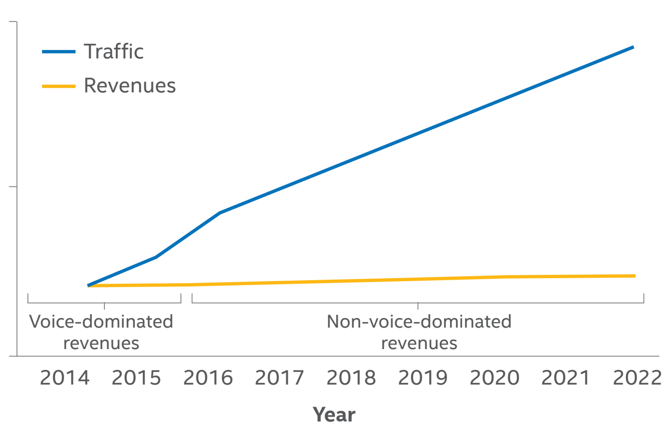

Figure 1. Traffic growth[3] vs revenue trends[4]

In addition to the downward revenue pressure, CoSPs are also facing increased competitive challenges from OTT service providers for profitable internet-based services. The announcement from Google* on Project Fi*[5] signals their intent to enter the network service provisioning market. Facebook's* project TIP*[6] is a clear OTT drive to collaborate on new technologies, examine new business approaches, and spur new investments into the telecom space. The Google and RailTel* partnership and the Express Wi-Fi initiative by Facebook are enabling Internet access in areas lacking cellular connectivity. New entrants such as OneWeb*[7] plan to launch satellites in 2019 to provide high speed broadband across the globe. These companies develop their solutions and OTT-hosted services in software on cloud infrastructure and provide rapid service releases and updates to their customers, and in this model the network is merely a delivery pipe to their target customers. In response, CoSPs must redefine and transform how they deliver new and existing infrastructure and services to enable them to compete with OTTs in this new order.

The adoption of software programmable networks offers CoSPs the capability to be more strategic about when to partner and when to compete with these players. It also allows CoSPs to simultaneously transform their internal organizations, procedures, Operation Support Systems (OSS), and customer engagement models. Thus CoSPs can bring the required automation (for cost reduction and efficiency) and service agility necessary in a digital services world.

1.2 Evolution to Software-centric, Distributed, Converged Network

Since the 2012 publication by ETSI of the NFV white paper, the industry's standards activities having been moving toward making the promise of the technology a reality. Currently, there are multiple ongoing standardization efforts at the Broadband Forum for fixed networks, at 3GPP for wireless networks, and at other consortia looking at various aspects of an NFV/SDN-centric network architecture such as CORD. These standards demonstrate the maturity of the problem statement and willingness for CoSPs and vendors to transform the CoSP's edge and central office environments to enable these new flexible service delivery paradigms.

Figure 2 illustrates some of the key virtualized components that make up a distributed architecture. The level of distribution will vary depending on the services supported by the network. We expect to see mobile infrastructure distribute to mid-mile sites in order to serve the massive and sustained traffic CAGR. For very low latency applications, the mobile user plane will distribute right out to the last mile site or even close to the base station itself. Another emerging transformational force is the ongoing upgrade of the broadband wireline infrastructure, in 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"—that, is 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 CO facilities and thus we need many fewer distributed COs to serve a given population.

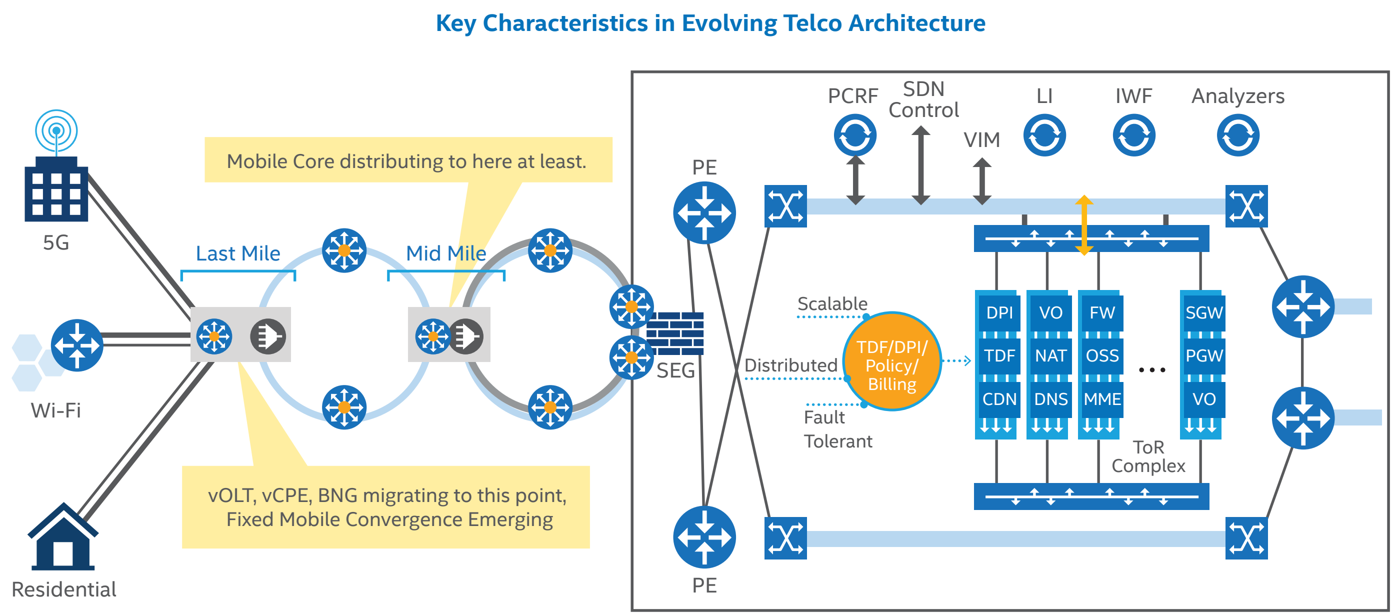

Figure 2. Evolution of telecom architecture

These new COs also house new broadband radio technologies in order to initiate fixed wireless access (FWA) bearers to serve the last mile complementing 5G and PON upgrades.

On top of these initiatives, CUPS, which supports separate virtualization of a VNF's data plane and user plane, and CORD are also actively being pursued by many CoSPs and will change how nodal functions may be designed, and how sites are built and orchestrated. These efforts amount to a network infrastructure revolution with many moving parts.

Fixed network CoSPs are adopting disaggregated networks where the network functions (e.g., firewall, BNG, SBC) are independent from the hardware that implements the transport packet forwarding in the data plane.

A similar service disaggregation is happening in wireless networks driven by 5G's main use cases for enhanced broadband, ultra-low latency and reliable communications for critical applications, as well as scalable machine-to-machine (M2M) communication. The backhaul requirements for 5G also require extensive infrastructure with substantial quality of service (QoS) requirements in terms of throughput, latency, jitter, and synchronization.

This all leads to a convergence happening in the edge of the network, whereby CoSPs across all geographies are exploring fixed-mobile convergence infrastructure (FMC) at the network edge and in the last mile. There is a scramble for relevance in this new architecture across CoSPs, TEMs, and silicon manufacturers. The concept of the next generation central office (NGCO) has now emerged, which is a fiber-rich central office that will support both fixed and mobile operators and be capable of serving more subscribers. It will implement functions in a more software-centric way, which will allow it to deploy new and more flexible services. It is intended to function as a local edge data center with a smaller area and power footprint than the traditional centralized data center.

The Intel NGCO mini data center architecture is shown in Figure 3 below. This comprises industry standard servers and standard storage and switching infrastructure. As part of a reference NGCO design, Intel will enable several gateway types (EPC, BNG, and CMTS) within the same complex—thus demonstrating the delivery of FMC requirements for NGCO, based on a flexible mix of bare metal, virtual machines, and containers. The mini data center will also demonstrate associated service delivery—again using industry standard servers.

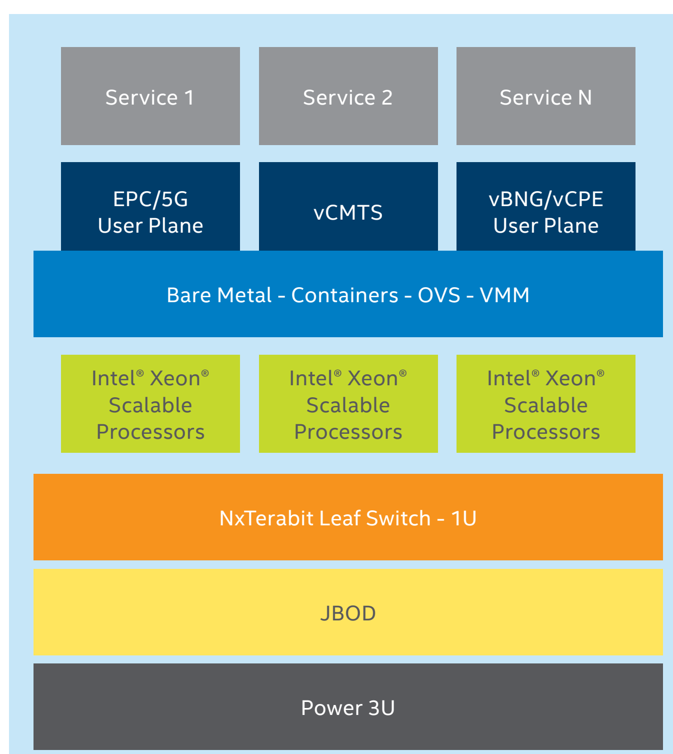

Figure 3. Intel architecture for NGCO mini data center

1.3 Evolution to Programmable Networks

With the advancement of the new network architecture, traditional network equipment providers are adopting more configurable, more programmable virtualized architectures that can support CUPS, SDN, and new services as required.

In parallel, new highly programmable solutions are emerging based on field programmable gate arrays (FPGAs), network processing units (NPUs), or other proprietary programmable ASICs, which boast support for networking programming languages such as P4. P4 (programming protocol-independent packet processors) is an open source domain packet forwarding programming language supported by many silicon vendors.[8] Such languages promise greater flexibility than traditional ASIC appliances and offer the ability to program new user plane encapsulations, metadata, and behaviors without additional silicon development cycles, albeit still at the cost of having to compile the network application for each required technology.

However, the technology that has seen the biggest growth and adoption rate is network functions virtualization enabled by software running on industry standard servers. NFV solutions provide extreme flexibility, scalability, and code portability (ecosystem), and their performance has been steadily improving (as outlined in Section 3.3) to the point where these servers have the throughput that makes them ideal for edge NGCO deployments.

2 Evolved Network Architecture

2.1 Benefits of Software-centric Networking

As the industry gains more experience with NFV, CoSPs are seeing that the flexibility afforded by a software-centric approach enables a single edge NFVI platform to serve the emerging, and often divergent, service needs of consumer, residential, and enterprise customer segments in both fixed and wireless networks. This software-centric approach also allows CoSPs to quickly deploy new services, including in localized rollouts for market testing where the cost of failure and subsequent service removal is minimized.

2.1.1 Flexible Service Provisioning

Distributed networks provide a much more flexible approach to service provisioning. For instance, as the number of IoT devices increases, they provide a much larger target footprint for security attacks. Large distributed denial of services (DDoS) attacks can generate excessive traffic and threaten the network stability, as recent DDoS attacks have shown.[9] Preventing such attacks at the edge has enormous benefits because the attack is blocked at its origin, thus avoiding the load and disruption to core or off-net services. Distributed NFV security applications are ideally suited for edge security. For example, DDoS blocking filters can be configured and deployed to neutralize the attack within seconds. Such an approach can also simplify the security of new applications and devices by shifting the security implementation to an edge NFV.[10]

For enterprise customers, the vCPE use case was envisioned as a platform with the compute power and remote deployment flexibility that allows CoSPs to sell optional services (firewall, intrusion prevention, WAN acceleration, unified communications, etc.) above core networking services. These services are computationally intensive and therefore very amenable to implementation as VNFs in software. This approach enables a much less costly “trial before buy” approach to managed service upselling, making it significantly cheaper to deploy on-net service trial bundles and removing the need for expensive CPE upgrades and site visits for new services, while improving time to revenue possibilities.

In mobile networks, vEPC/v5GCN (virtualized evolved packet core/virtualized 5G core network) will facilitate the provisioning of additional functions and services at the edge and enable seamless transition of the network to 5G through the remote upgrading of these VNFs. Very low latency applications such as AR/VR, tactile Internet, automation, and connected car can then be provisioned on demand where and when needed.

Another considerable advantage of dynamic service provisioning at the edge is the possibility to deploy and prioritize the most needed services locally during periods of network disruption, thus increasing the resilience of the network. This is especially important in disaster zones where infrastructure can be damaged. The inherent resilience of distributed NFV networks was a key catalyst for the NFV initiative.

Provisioning flexibility also allows the provider to provision the same NFVI resources with different needs depending on time of day: for example, dynamically dimensioning between BNG (broadband network gateway) and EPC gateways as customers move their data usage from their homes in the suburb (broadband) to their work in the city (mobile) and home again over the course of the day. The suburban BNG resource is expanded through the remote instantiation of new VNFs in the evening as homes are streaming video, and reallocated to EPC during the day when the fixed suburban network is at lower load.

2.1.2 Fast Service Provisioning

SDN/NFV also helps to close the loop between the customer ordering a service and when that service is provisioned. CoSPs are currently transforming their OSS and BSS systems to give their customers instant ordering portals through the integration of SDN controller and NFV management and orchestration (MANO) stacks where these new services are instantly created in software. This level of automation and adaption allows the CoSPs to transform service provisioning from a mostly manual process that involves stocking or ordering specific hardware and configuring OSS controllers to a process where the service provision is much more automated and remote.

Today the provision of consumer (data, voice, and video) and enterprise (e.g., VPN) services in the fixed network typically involves the configuration of different types of service functions on a multi-service edge router. The BNG typically provides residential broadband Internet services while the provider edge router (PE) typically connects enterprise estates, providing them with global VPN connectivity through the operator's multi-protocol label switching (MPLS) network.

The move to edge-hosted VNFs gives the operator a more flexible approach to the provisioning of such services. For example, rural non-commercial regions require only BNG deployment, whereas in urban or commercial centers, the BNG is needed for urban dwellers and the additional PE functions are required to provide the requisite enterprise services. It is advantageous for both cost and time to deploy virtualized BNG and PE services on industry standard servers, which have the performance range suitable for edge deployment while providing extreme flexibility and the possibility to reuse hardware for multiple services. The same BNG or PE router VNF can run on any standard Intel architecture platform in any location, thus enabling an extremely portable and scalable architecture.

Similarly in mobile networks, consumer and enterprise services typically pass through different network functions. Consumer traffic passes through Gi-LAN functions such as deep packet inspection (DPI), video/web optimization, and CGNAT, whereas enterprise traffic typically passes through VPN functions. These VNFs are deployed on an as-needed basis as per the fixed domain.

Forthcoming 5G networks will be required to handle high data throughput and low latency, thus necessitating a distributed network architecture and edge hosting of some core network functions. Edge-deployed VNFs are best placed to handle the flexibility and dynamic service provisioning required for mobile users. It makes sense to share resources between enterprise and consumer traffic and dynamically re-allocate them depending on the demand of the particular location and time of day, thus reducing the required overprovisioning compared to using dedicated appliances.

As an evolution of this, the same Intel architecture resources may be dynamically shared between VNFs serving fixed, residential, and enterprise services, and we see 3GPP and Broadband Forum (BBF) beginning technical studies for a similar standardization now.

2.1.3 New IT-Based Network Cost Models

Adopting software-based network services opens up opportunities for CoSPs to pursue usage-based or "follow the sun" enterprise licensing models. In this model, instead of paying for their services and equipment expansion in a traditional CapEx-driven budgeting model, CoSPs can move to a more usage-based OpEx model, which is better aligned with revenue projections. NFV/SDN is re-forming the vendor ecosystem by allowing CoSPs to decouple software from hardware, thus achieving a much more optimized supply chain. The increased competition is leading vendors to be more innovative and flexible in the type of licensing they offer for their networking VNFs.

New software licensing models have replaced the traditional perpetual licenses with included upgrades with more granular, pay-per-upgrade or subscription-based license models that may even include hardware lease to provide network functions as a service. The unit of measure for license sizing is also evolving from licensing based on installed networking capacity to licensing based on actual use, number of users, transactions, or even use-case-specific billing.

Such an approach provides enormous flexibility for CoSPs to pick and choose the most cost-efficient licensing model and indeed a supplier that best meets the required service requirements.

Flexible licensing also directly benefits enterprises and consumers as they can pick and choose the most appropriate mix of services and functions and their desired function and price point.

2.2 Control and User Plane Separation

NFV more easily enables a migration to a CUPS architecture where the control plane can be deployed at a different network location from the corresponding user plane implementation. 5G architecture will rely extensively on CUPS to achieve independently scalable control and data planes. EPC equipment is already CUPS ready from 3GPP Release 14 and allows a mix and match between hardware user plane and virtualized control plane. CUPS standardization in fixed networks[11] is in progress, and once in place it will enable virtualizing those network functions where this is beneficial. For example, a virtualized control plane can be elastically scaled up in periods of high load, such as reconnects after outage, independent of the constraints of a hardware-based user plane implementation.

The disaggregation of services from the packet forwarding engine in mobile networks is illustrated in Figure 4 and for fixed networks in Figure 5 respectively. In a VNF implementation, the network functions are independent from the hardware that implements the packet forwarding in the data plane and can run anywhere on any standard server or in the cloud.

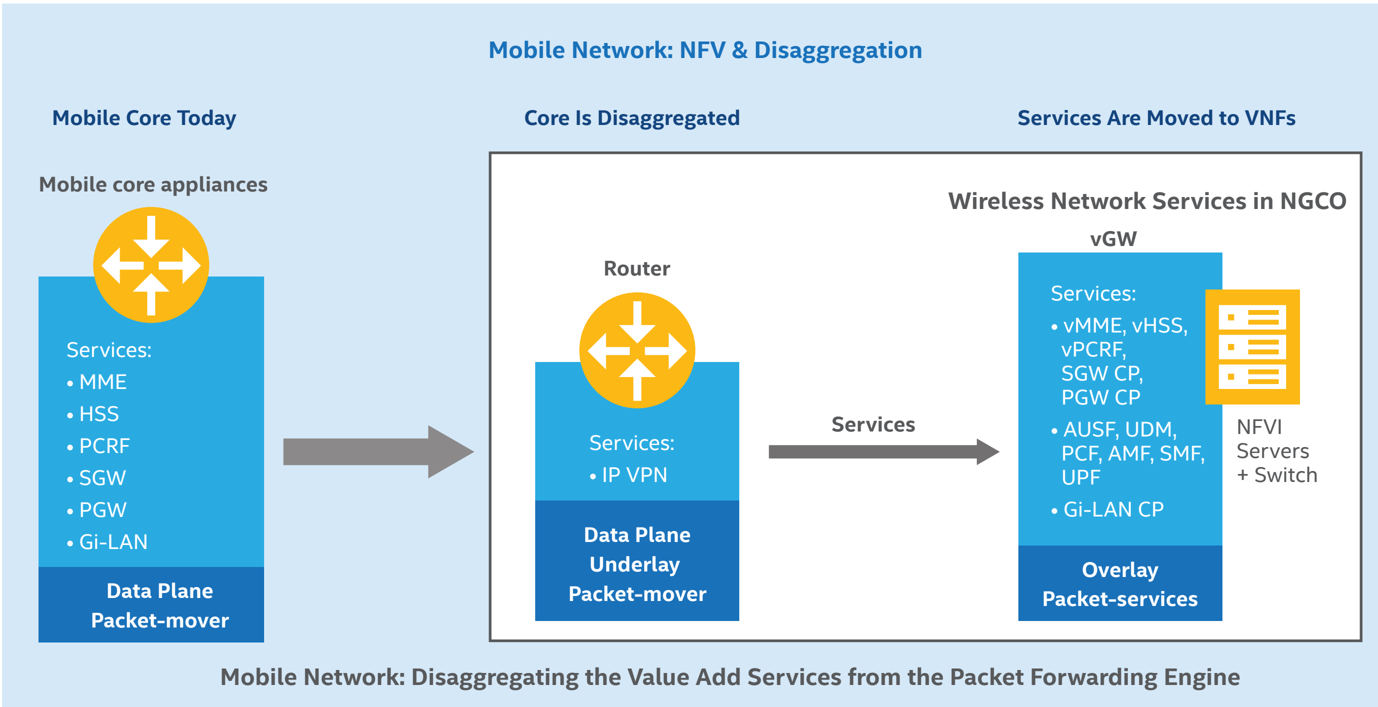

Figure 4. Service disaggregation in 4G and 5G mobile networks

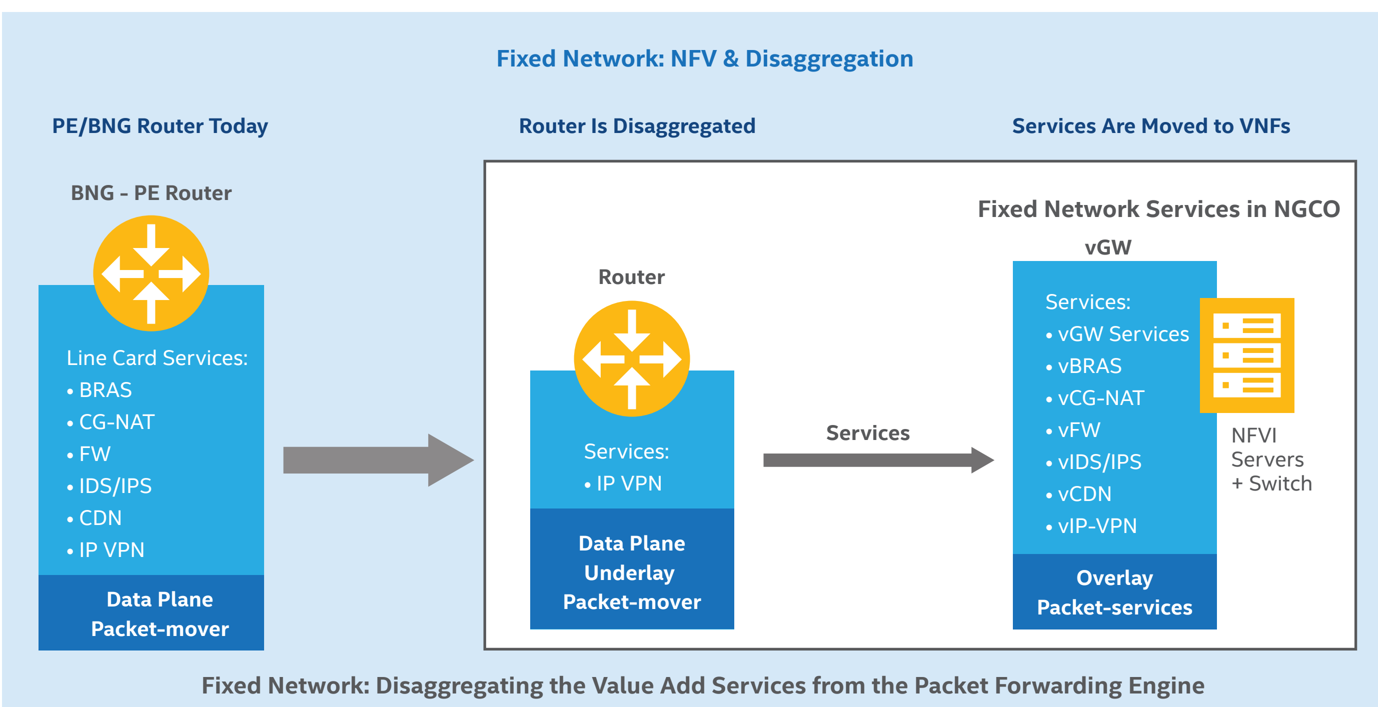

Figure 5. Service disaggregation in fixed networks

2.3 Edge Deployment

NFV allows operators not only to decouple network functions from hardware, but also to decouple network functions from location. Thus it can place functions in a cost-effective way where they are required, including at the NGCO or indeed deeper inside the operator network if required.

2.3.1 Network Service Requirements

Improvements in network reliability are an important aspect of edge network deployments. Today's typical equipment and service uptime requirements mandate hardware redundancy models, which requires replicating appliances or blades in a 1:1 or 1:N manner to ensure all subscriber traffic is protected. The protection (software and hardware) device must be supplied by the same vendor. This is expensive from the perspective of acquisition, real estate, and maintenance. NFV enables more flexible N:1 or M:N application redundancy models where a VNF (e.g., vEPC) can be instantiated as a standby software entity hosted on industry standard servers in order to protect many more active VNF instances at a lower cost protection point per subscriber. This advanced redundancy can also be implemented in a hierarchical model whereby a centralized protection VNF may protect many access-located VNFs. In the case of a failure, the traffic can quickly be switched to the standby VNF instance. This is illustrated in Figure 6.

A flexible edge/NGCO platform approach enables CoSPs to adopt SDN/NFV best practices to build the scalable edge enterprise, mobile, and residential services as envisioned at the time of the ETSI ISG formation. It enables the deployment of all of the software VNFs needed for user plane (vCPE, vBNG, vEPC), advanced enterprise security (virtual firewall - vFW, intrusion prevention system - IPS), consumer security (parental control) and video services (vCDN, virtual set-top box - vSTB) to be instantiated on the same edge platform.

The flexibility of the approach opens up the opportunity for revenue sharing models with OTT/service vendors that need access to such critical network-enabled infrastructure to improve the latency or geographic location of their own services.

2.3.2 Technology Drivers

With the consumer-driven move to OTT service consumption, the impending arrival of 4K video combined with the growing number of connected home/enterprise devices, and the expected ramp in machine-to-machine (M2M) device connectivity will lead to sustained increases in data traffic on the order of 30 percent CAGR across fixed and wireless networks.[12] To deal with this immense traffic increase, CoSPs are planning to move traditionally centralized functions to the edge for both fixed and wireless applications. This reduces the load on core networks, enables the CoSP to distribute content, and allows for new services based on the distributed architecture.

For example in mobile, EPC user plane is already migrating to the metro-area networks, as seen in Figure 8 (Distributed edge data plane and services). Fixed wireless access will exacerbate this trend further because of the large volumes of traffic that need to be supported. Also, as the virtualization of the radio access network (vRAN) is becoming more feasible, vRAN baseband units are also being deployed in the last mile.

Similarly in fixed networks, large-scale FTTx PON deployments (vBNG) and the rising popularity of virtual optical line termination (vOLT) and vCPE solutions are driving the migration of these functions to the last mile location.

2.3.3 NGCO and the Opportunity for Collaboration

AT&T* and Amazon* have recently announced an expansion of their partnership to enable an easier and more secure interconnect between Amazon's cloud services and AT&T FlexWare* edge platform, with the aim to help businesses increase the agility of their edge computing deployments.

In addition, AT&T and Amazon are exploring options to bring Amazon's Greengrass IoT* platform to AT&T FlexWare and open new IoT opportunities for businesses. It is not difficult to envisage a scenario where CoSP-owned edge NFVI delivers the long-awaited two-sided revenue model driving OTT collaboration with APIs provided from locational-relevant CoSP operations in areas as diverse as edge analytics, third-party CDN, government, public safety, private security, and gaming.

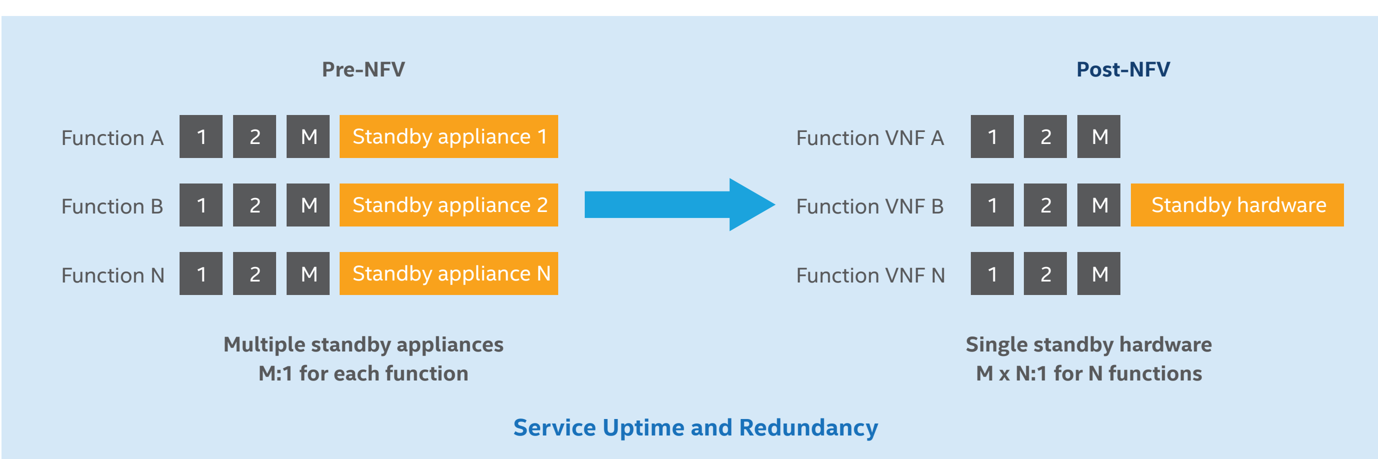

Figure 6. Service uptime and NFV-SW redundancy

Most importantly this also enables the CoSPs to move from a model where these services are tightly coupled to traditional vendors' hardware and OSS controllers to one where the service provision is much less proprietary and is open to a larger swath of independent and third-party software service vendors—providing greater options for innovation.

2.4 5G and Fixed Mobile Convergence

CoSPs want to go beyond hosting separate network functions on a common infrastructure and achieve truly converged fixed and mobile networks. FMC initiatives are driven by the necessity to protect CoSP network investments and revenues. This is becoming increasingly critical to CoSPs as we approach the 5G spending wave. To achieve this, CoSPs aspire to have common credentials, policy, and user data management between fixed and mobile networks. This will ultimately result in converged control and user planes between EPC for mobile and BNG for fixed networks.

3GPP and the Broadband Forum already have a joint initiative underway with aligned time frames to achieve FMC standardization, with study work planned in 3GPP Release 15, and normative work in Release 16.[13] The high-level architecture is shown in Figure 7.

A software-centric network based on NFV/SDN will enable the sharing of resources and avoid duplication of network functions and separate over-provisioning in fixed and wireless networks. The emerging CUPS architecture allows convergence at any level of the network.

However, the convergence is most beneficial in the network edge, where the performance requirements match the capabilities of industry standard server-based NFV solutions and enable full virtualization for an extremely flexible architecture. The evolution to a converged fixed/5G smart edge where the user plane is distributed to enable services is summarized in Figure 8.

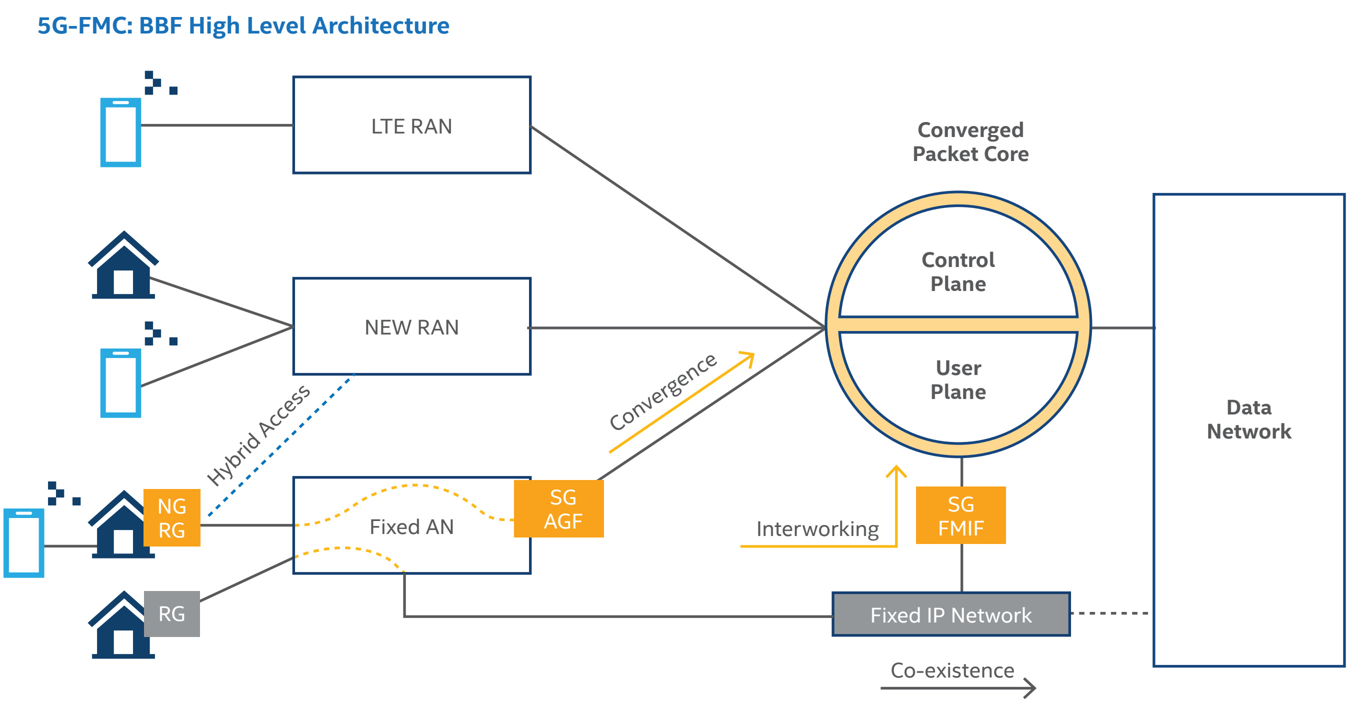

Figure 7. Proposed BBF/3GPP FMC architecture

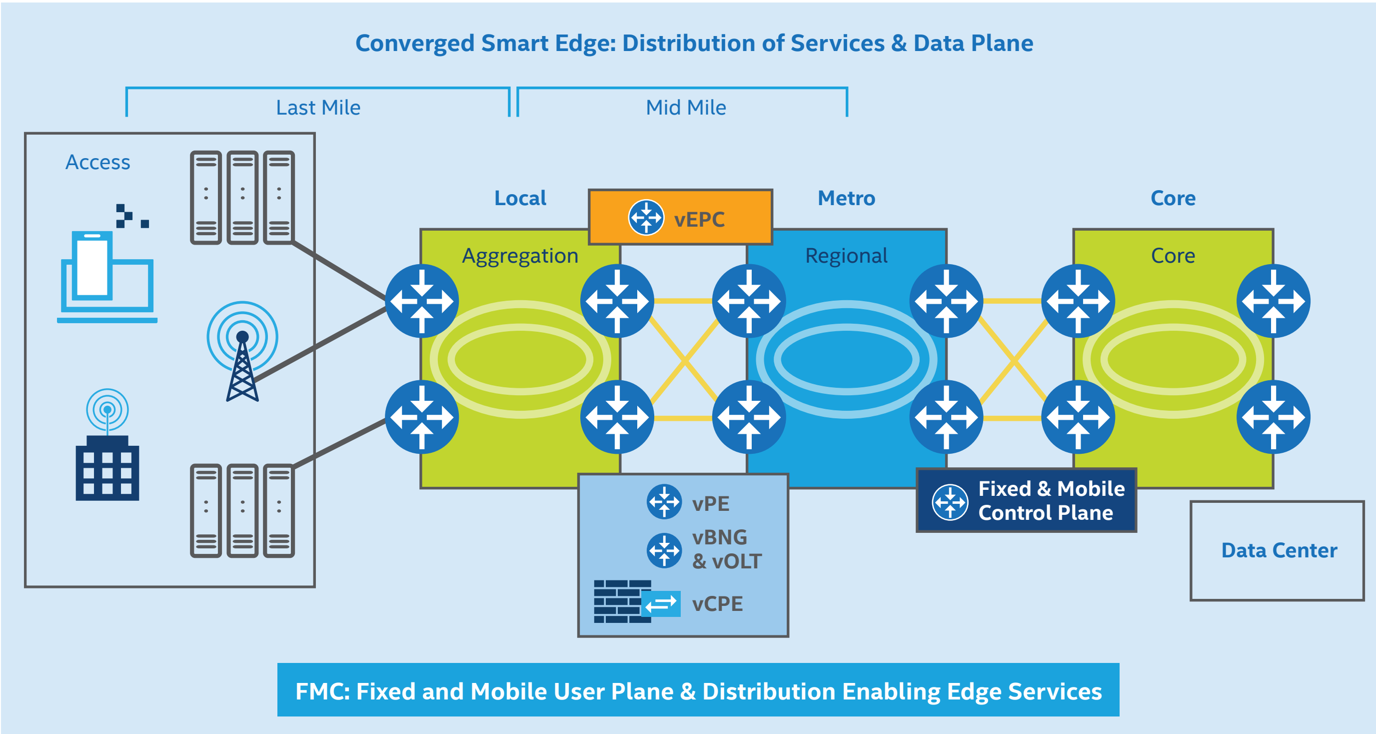

Figure 8. Distributed edge data plane and services

3 Networking Technologies

3.1 Traditional Networking Platforms

Purpose-built hardware networking platforms based on proprietary ASICs or NPUs are usually performance leaders in the user plane, currently reaching n x Tbit of traffic throughput per line card instantiation. Interim software releases provide (mostly) control plane functionality improvements on a six-to-nine-month software release cycle. However, the upgrade capabilities are constrained by the hardware platform because there is such a tight coupling between software and hardware. Service innovation is directly tied to the hardware development cycle, on the order of 24 to 36 months, which contributes to the slow pace in introducing new services. In the current commercial environment, where CoSP competition has shifted from like-minded CoSPs to more nimble, Internet-speed OTT providers, this approach to service innovation places them at a significant disadvantage.

Traditional ASICs are expensive to design, test and manufacture. The required development skill sets—e.g., register-transfer level (RTL) design and silicon validation—are expensive and in short supply. Similarly, in software development there is a limited pool of developers with deep knowledge of the proprietary platform, which makes it more expensive and time consuming to develop software. Then these upfront development and manufacturing costs have to be recouped across the product life cycle and can become a substantial price driver for products with low sale volumes. The huge cost of failure makes service innovation slow and expensive.

This is not to say that ASICs do not have their place in next generation networks. ASICs may always be used in the fixed-function transport domain where the transport protocol stack (layer 1-4 OSI) is well understood/standardized and aggregate switching speed is the premium technology choice factor. However, and precisely because of, the new competitive situation and need for service agility, enterprise and consumer-oriented services (L5-7) must now evolve at a much faster pace, and these phenomena drive their disaggregation from the main transport/switch function to more suitable software-based implementation.

The P4 initiative[8] aims to address that trend and develop a unified network programming language that can be compiled on a variety of platforms, including FPGA and NPU-based appliances. However, P4 code still needs to be compiled individually for each hardware platform, and any compiler compatibility issues, which are inevitably present between platforms even in the most mature programming languages, will limit its portability. Also, the need for different hardware-specific images will limit the elasticity of software instances and make resource management more complicated.

In addition, the need for application-aware security, especially as the security perimeter distributes to the NGCO locations, leans toward Intel architecture being optimal for multi-service platform deployments at those locations as (a) they are more than performant enough for those locations, (b) they are truly programmable, (c) they can implement security policy uniformly and consistently from core to edge, and (d) it enables a truly hardware independent ecosystem of software providers to emerge.

3.2 Intel Architecture's Networking Platform

The performance of virtualized networking functions has been steadily improving as NFV has matured. Virtualized software is becoming more advanced and efficient as the understanding of NFV evolves. And hardware performance is also improving not only through the usual performance improvement cycles of Moore's law, but also because of a range of network hardware acceleration features that have been added to industry standard servers. The performance of VNFs has now reached a level that makes them an attractive option to deploy in sites with small to medium loads, such as NGCO edge sites.

The advantages of VNFs are numerous. The software is completely portable between different servers and allows a very elastic deployment with fast ramp-up or -down between different deployment instances. The upgrade cycles are short, typically three to six months, allowing continuous upgrade of services and capabilities. It also allows a full decoupling between the software and hardware suppliers, thus reducing the dependency on a single source and increasing the solution flexibility at a more favorable price point.

In addition, NFV platforms allow a more efficient handling of asymmetric traffic as the platform resources can be allocated dynamically depending on the uplink/downlink traffic requirements, as shown in Figure 9. In both wireline and wireless traffic, there is asymmetry between uplink/upstream and downlink/downstream traffic. For example, traffic statistics from the mobile networks in Japan show 1:6.6 traffic ratio between uplink and downlink.[14] The ratio is similar for consumer traffic in fixed networks; only enterprise traffic is close to symmetric. In traditional ASIC/NPU-based appliances, very often the allocated silicon resources are fixed in terms of uplink/downlink throughput. In a virtualized solution, the silicon resources, such as CPU cores, RAM, etc., can be assigned and re-assigned to traffic functions in either direction as required.

This bandwidth and resource agility can also help where new applications are being trialed, such as emerging augmented reality/virtual reality (AR/VR) services, where the possibility of sudden adoption/uptake could rapidly alter the upstream/downstream ratio and traffic throughput in a given region.

3.3 Silicon Applicability

In telecommunication networks, appropriate silicon selection for the platform depends on the application, throughput, and deployment location, as illustrated in Figure 10. Centralized core deployments require very high throughput but can tolerate lower flexibility if this results in economies of scale, and ASIC/NPU-based transport fabrics may be preferred for such locations. Edge locations, on the contrary, need to support lower throughputs but require extreme flexibility in running a multitude of services on the smallest possible number and variety of networking equipment. NFV based on Intel architecture is the clear choice for edge locations.

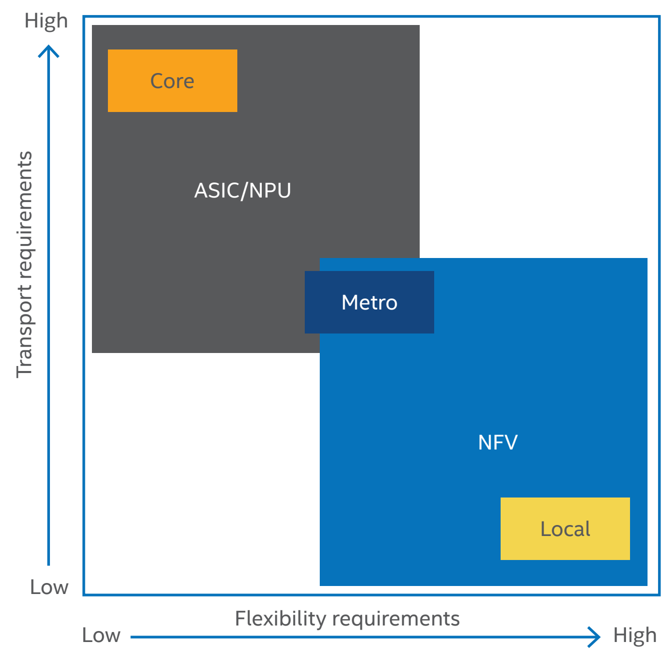

Figure 10. Transport and flexibility requirements of raw transport fabric

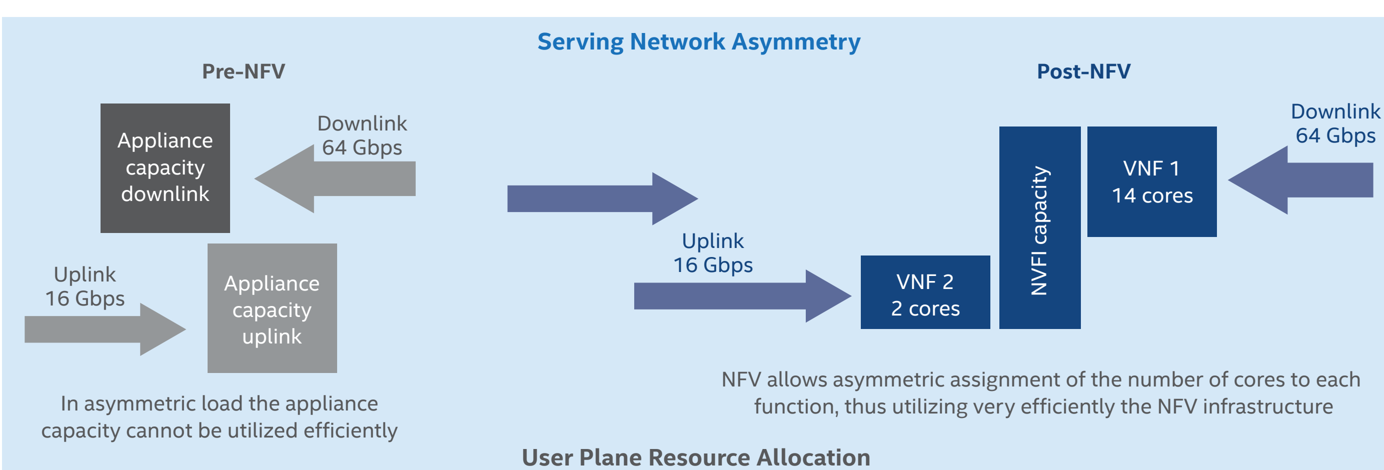

Figure 9. User plane resource allocation

Existing last mile edge locations typically terminate between several hundred to thousands of subscribers[15] in EMEA/US, with higher numbers in the densely populated urban centers in the APAC regions. However, the copper melt, i.e., the replacement of legacy copper access COs with optical fiber-served COs, is continuing at a fast pace and is redefining the scale of the last mile NGCO. The longer fiber reach means last mile will be aggregated into fewer sites, and as a result the local NGCO will serve a larger number of customers, possibly reaching to tens of thousands. For metro sites, the number of served subscribers is on the order of hundreds of thousands of subscribers. Both cases serve lower traffic volumes than in centralized deployment approaches. A summary of typical subscriber volumes and corresponding expected traffic volumes in 2017 for fixed and wireless networks at each network location is summarized in Figure 11.

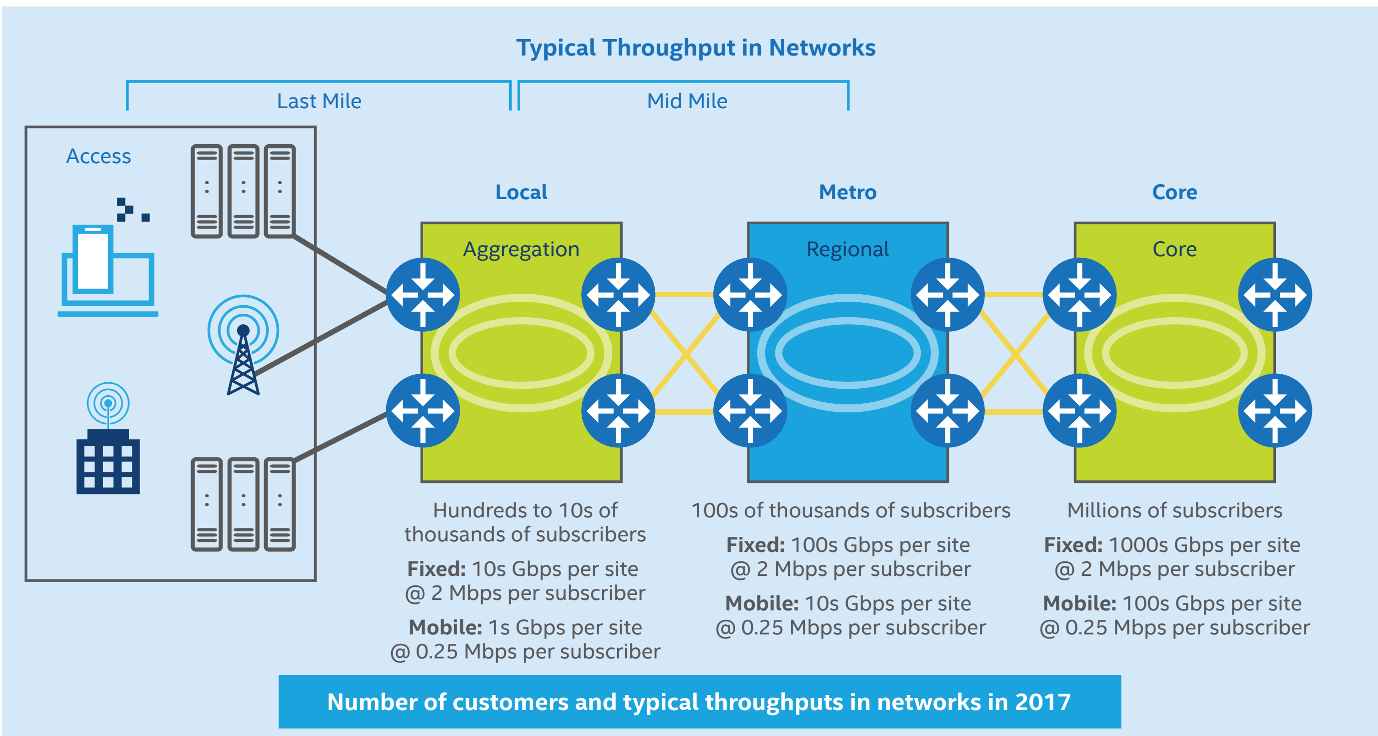

Figure 11. Number of customers and throughput in networks in 2017[16]

Advances and new techniques have delivered a substantial improvement in the performance of NFV solutions. Performance tests demonstrate that NFV is already capable of serving last mile and mid-mile traffic volumes for both fixed and wireless network functions. A comparison between an actual vBNG solution and the edge BNG requirements is illustrated in Figure 12, whereas a comparison between a vEPC solution and the edge EPC requirements is shown in Figure 13. In both cases, the NFV solution demonstrates performance levels that match and exceed the requirements of last and mid-mile deployments on a single server, while providing much needed flexibility to also run other network functions on the same hardware.

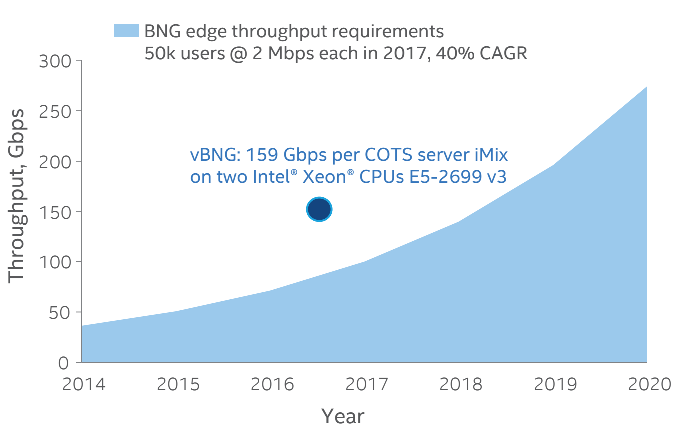

Figure 12. Edge BNG requirements vs NFV performance[17]

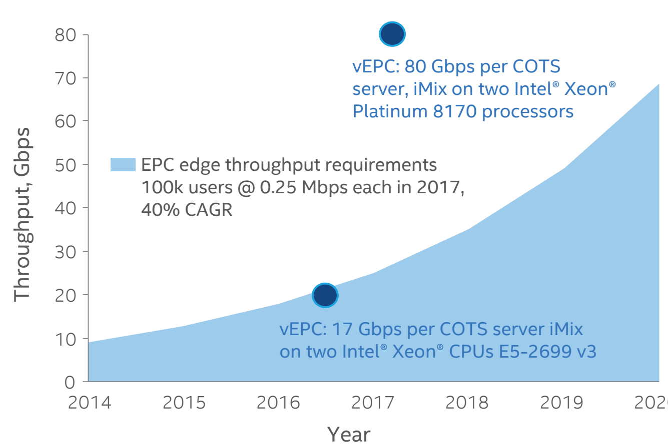

Figure 13. Edge EPC requirements vs NFV performance[18]

3.4 Advanced Intel Architecture Networking Technologies for NGCOs

The constant addition and expansion of advanced technologies in Intel architecture and supporting technologies allows the rapid performance improvement of NFV solutions. Some of the technologies include:

- Hardware acceleration of load balancing functions using techniques such as the Data Plane Development Kit (DPDK), poll mode driver (PMD), and dynamic device personalization (DDP), which can improve significantly the performance in comparison with a software load balancing solution.
- Implementing a run-to-completion (RTC) model for VNF packet processing, thus aiding the real-time processing capabilities essential for networking.
- Use of hardware-accelerated packet forwarding, such as SR-IOV.
- Balanced I/O systems, which allow faster discovery, provisioning, and improved performance of peripherals, and ensure deterministic behavior and more efficient use of two-socket servers.
- Use of hardware-accelerated encryption/decryption functions built into Intel® QuickAssist Technology (Intel® QAT).
- Inclusion of FPGA on network interface cards, which allows vendors to implement innovative targeted hardware acceleration techniques to improve performance.

All these techniques have allowed NFV solutions based on Intel architecture to achieve performance levels that put them in the sweet spot for edge deployment of networking functions.

4 Conclusions

CoSPs have been driving a major network transformation to achieve greater efficiencies, flexibility, resilience, reduced costs, and new business opportunities. Several major trends are evolving within this transformation:

- NFV/SDN, which moves networking towards software-centric solutions and allows extremely flexible dynamic service provisioning while decoupling software from hardware and making the supply chain leaner and more efficient.
- CORD, which reflects a major drive to deploy functions to the edge to support new use cases and business models
- CUPS, which decouples the control from the user plane, thus allowing deployment of each plane individually in the most efficient configuration and location.
- 5G, which will enable many new mobile use cases and open new opportunities for mobile network operators.
- FMC, which will allow full convergence of control and user plane between mobile and fixed networks and enable CoSPs to protect their investments and revenues with new user services.

While traditional ASIC and NPU-based networking platforms can provide very high throughput, they lack the flexibility and portability essential for agile edge deployments.

Intel has introduced a number of new technologies in Intel architecture to enable the NFV performance improvements, which are continuing at a high pace. As a result, NFV performance has been steadily increasing and throughput performance has reached levels where industry standard servers can comfortably support NGCO edge deployments of network functions in a cost-effective way.

NFV has become an extremely compelling choice for edge deployment, as it also gives added flexibility and portability, and allows economies of scale, new functions, and new business opportunities that were unthinkable in hardware-centric solutions.

Glossary

Term	Description
AR/VR	Augmented Reality / Virtual Reality
ASIC	Application-Specific Integrated Circuit
BNG	Broadband Network Gateway
CAGR	Compound Annual Growth Rate
CDN	Content Delivery/Distribution Network
CGNAT	Carrier-Grade NAT
CO	Central Office
CORD	Central Office Re-architected as a Datacenter
COTS	Commercial Off-The-Shelf
CPU	Central Processing Unit
CoSP	Communication Service Provider
CUPS	Control User Plane Separation
DDoS	Distributed Denial of Service
DDP	Dynamic Device Personalization
DPDK	Data Plane Development Kit
DPI	Deep Packet Inspection
EPC	Evolved Packet Core
ETSI	European Telecommunications Standards Institute
FMC	Fixed Mobile Convergence
FPGA	Field Programmable Gate Array
FW	Firewall
FWA	Fixed Wireless Access
IPS	Intrusion Prevention System
ISG	Industry Specification Group
MPLS	Multi-Protocol Label Switching
NAT	Network Address Translation
NFV	Network Functions Virtualization
NFVI	NFV Infrastructure
NIC	Network Interface Controller
NPU	Network Processing Unit
PE	Provider Edge router
PMD	Pole Mode Driver
OLT	Optical Line Termination
OSS	Operations Support Systems
OTT	Over-The-Top
RAN	Radio Access Network
RSS	Receive Side Scaling
RTC	Run To Completion
RTL	Register-Transfer Level
SDN	Software Defined Networking
SR-IOV	Single Root I/O Virtualization
STB	Set-Top Box
TEM	Telecom Equipment Manufacturers
VIM	Virtual Infrastructure Manager
VM	Virtual Machine
VNF	Virtual Network Function

Platform References

Reference platform from Light Reading* EANTC*-Nokia* test report

Datapath VM

Processor (Type, clock speed)	Intel® Xeon® CPU E5-2699 v3 @ 2.30GHz (18 core)
Memory in KB	65934776
NICs	4x Intel® Ethernet Server Adapter X520-2 (82599 chipset)
PCle* version	PCie 2.0 x8

Control VM

Processor (Type, clock speed)	Intel Xeon CPU E5-2687W v3 @ 3.10GHz (10 core)
Memory in KB	131998316
NICs	4x Intel Ethernet Server Adapter X520-2 (82599 chipset)
PCle version	PCle 2.0 x8

SAEGW (OAM, LB and MG VMs)

Hardware	HP* C7000 Blade system
Blade Type	ProLiant* BL460c Gen9 Server Blade
Interconnect Bay	6125XLG Blade Switch
Processor (Type, clock speed)	2x Intel Xeon CPU E5-2680 v3 @ 2.50GHz (12 core)
Memory in GB	128
NICs	HP Ethernet 10Gb 2-port 560FLB
PCI Mezzanine Card	HP Ethernet 10Gb 2-port 560M

MG for ePDG

Processor (Type, clock speed)	Intel Xeon CPU E5-2699 v3 @ 2.30GHz (18 core)
Memory in KB	65934776
NICs	4x Intel Corporation Ethernet Server Adapter X520-2 (82599 chipset)
PCle version	PCle 2.0 x8

LB and OAM for ePDG

Processor (Type, clock speed)	Intel Xeon CPU E5-2687W v3 @ 3.10GHz (10 core)
Memory in KB	131998316
NICs	4x Intel Corporation Ethernet Server Adapter X520-2 (82599 chipset)
PCle version	PCle 2.0 x8

Software Running on Host

Host operating system and kernel version	CentOS* Linux* release 7.0.1406 (Core) 3.10.0-123.9.3.el7.x86_64
Libvirt version	libvirt-1.2.17-13.el7_2.2.x86_64
QEMU/KVM* version	qemu-kwn-ev-2.1.2-23.el7_1.8.1.x86_64

Cited Resources

[1] https://portal.etsi.org/NFV/NFV_White_Paper.pdf

[2] See IHS Telecom Trends and Drivers H1 '17, p. 13: https://technology.ihs.com/589822/telecom-trends-drivers-market-report-regional-h1-2017

[3] http://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.pdf

[4] https://www.gsma.com/

[5] https://fi.google.com/about/

[6] https://telecominfraproject.com/

[7] http://oneweb.world/

[8] https://p4.org/

[9] See https://blogs.akamai.com/2016/10/620-gbps-attack-post-mortem.html and https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/

[10] http://ieeexplore.ieee.org/document/7945849/

[11] https://www.broadband-forum.org/standards-and-software/major-projects/cloud-central-office

[12] See Cisco VNI 2015-2026 http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.pdf, http://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.pdf, and http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html#MeasuringMobileIoT.

[13] http://www.3gpp.org/news-events/3gpp-news/1287-3gpp-and-the-broadband-forum-collaborate-on-fixed-mobile-convergence-standards

[14] https://www.itu.int/dms_pub/itu-r/opb/rep/R-REP-M.2370-2015-PDF-E.pdf

[15] https://www.fcc.gov/internet-access-services-reports

[16] Throughput and subscriber numbers for Figure 11 are based on interviews conducted by Intel with several tier-1 CoSPs across all major geos.

[17] 2017 BNG edge throughput requirements based on interviews conducted by Intel with several tier-1 CoSPs across all major geos, with 40 percent CAGR projection chosen as a high-case CAGR based on IHS Telecom Trends and Drivers H1 '17 research (see end note 1).

159 Gbps performance for a COTS server based on testing conducted by EANTC and Nokia, in which it was found that a single Intel Xeon processor E5-2699 v3 could achieve 79.5 Gbps throughput. Configurations: See Platform References section. Full report is available at http://www.lightreading.com/ethernet-ip/new-ip/validating-nokias-ip-routing-and-mobile-gateway-vnfs/d/d-id/720902.

[18] 2017 edge EPC requirements based on interviews conducted by Intel with several tier-1 CoSPs across all major geos, with 40 percent CAGR projection chosen as a high-case CAGR based on IHS Telecom Trends and Drivers H1 '17 research (see end note 1).

17 Gbps performance for a COTS server based on testing conducted by EANTC and Nokia, in which it was found that a single Intel Xeon processor E5-2699 v3 could achieve 8.5 Gbps throughput. Configurations: See Platform References section. Full report is available at http://www.lightreading.com/ethernet-ip/new-ip/validating-nokias-ip-routing-and-mobile-gateway-vnfs/d/d-id/720902.

80 Gbps performance based on testing conducted by Ericsson. Simulation of a radio access network with 2.6 million mobile broadband subscribers, 15 kbps average speed per user, and average 650 byte packet size. Configurations: 2x Intel® Xeon® Platinum 8170 processors, 2x Intel® Ethernet Converged Network Adapter x710 40 GbE. 2 EPG VMs run per server, each VM 50 vCPUs in size and equipped with 32 GB RAM. See http://cloudpages.ericsson.com/virtual-epc-capacity-evolution for more information.

Additional Resources

- http://www.etsi.org/deliver/etsi_gs/NFV/001_099/001/01.01.01_60/gs_NFV001v010101p.pdf
- http://www.3gpp.org/news-events/3gpp-news/1882-cups
- http://about.att.com/story/att_expands_relationship_we_aws.html
- https://arstechnica.com/information-technology/2016/02/netflix-finishes-its-massive-migration-to-the-amazon-cloud/
- Accelerating NFV Proto Application with Fortville Flexible Hash Filter by Andrey Chilikin and Xavier Simonart
- Data Center Solutions Intel® Ethernet Controller 700 Series Software Enhancements Dynamic Device Personalization (DDP): Tunnelled MPLS Use Case
- Intel Telecommunication Server Requirements Specification Document number 571467
- https://pages.nokia.com/16121.two.worlds.collide.LP.html
- https://www.itu.int/en/ITU-T/Workshops-and-Seminars/201612/Documents/slides/12-China_Mobile-FMC_network_architecture.pdf
- https://ark.intel.com/products/series/125191/Intel-Xeon-Scalable-Processors
- http://ark.intel.com/#@PanelLabel595
- https://01.org/intel-quickassist-technology