SOLUTION BRIEF

vEPC for Communications Service Providers

Executive Overview

Intel is accelerating adoption of network function virtualization with unique capabilities that enable optimal use of data center resources to deliver Communications Service Provider (CSP) services based on standard high volume servers. Virtualized Evolved Packet Core (vEPC) provides the CSP with the ability to reduce cost and drive new services by running previously hardware-dependent functions in a virtualized software model.

vEPC is being considered as a good means of evaluating the operational benefits of NFV technology in the Mobile Core, typically with Mobile Virtual Network Operator (MVNO) “slices” or Internet of Things (IoT) pilot projects, prior to adoption in Mission Critical Networks.

This document is aimed at senior managers evaluating NFV projects, and details Intel’s role in the ecosystem to enable CSPs to accelerate adoption of the technology and realize the business benefits of vEPC in production deployments.

Introduction

The exponential growth in mobile connections, traffic and workloads; the percentage of overall traffic from video; and the disparate types of devices driving the workloads have challenged the current Communications Service Provider (CSP) business model and require the CSPs to reconsider the approach and path to deliver, support and scale services. Network Function Virtualization (NFV) will provide a more cost-effective means to deliver mobile solutions and also enable development of more innovative revenue generating services for end-user consumption. Software-defined networking (SDN) and NFV will provide Communications Service Providers with the capabilities to revolutionize the industry.

The Evolved Packet Core (EPC) is the foundation of the mobile CPS’s core network through which services are created and delivered to end-users. Virtualizing the mobile core onto Intel® Architecture-based standard high volume servers (SHVS) will increase network flexibility, reduce costs, and enable CSPs to quickly launch new revenue-generating services more efficiently.

In the past, EPC solutions have been deployed on purpose-built systems. Today, CSPs are looking for alternative deployment models that will help them support fast-growing network demand. Moving to a virtualized EPC (vEPC) solution can help CSPs achieve more cost-effective scaling by using SHVS in place of purpose-built systems.

Implementing vEPC solutions can help CSPs scale to accommodate growth in the number of subscribers, traffic, or connections while controlling costs more effectively. Initially, CSPs might opt for small implementations of vEPC before applying this model to consumer voice and data traffic. Doing so allows them to explore the vEPC model and gain the necessary skills for managing these solutions before conducting a full implementation for consumer usage.

Many system integrators, software vendors, and platform providers rely on Intel® technologies to deliver the performance and scalability required for vEPC solutions. Using vEPC solutions based on Intel technologies, CSPs can offer a full range of services without having to use multiple purpose-built systems.

This paper describes the technologies required to develop and mature the virtualized evolved packet core that will enable the Communications Service Providers to alter the current state of the industry. Intel has the opportunity to work with the ecosystem to accelerate the adoption of the technology that enables CSPs to deliver virtualized mobile packet core production deployments.

Revision 1.3 April 2016

Table of Contents

Market Opportunity

The Mobile Packet Core is the foundation for CSP's IP-based service creation and enablement. The critical functions, which require high reliability from the mobile packet core, drive the rigidness of validation and operation constraints. The continued growth in the number of mobile connections and the amount and type of traffic has a direct correlation with the mobile network's capacity and scalability.

To maintain high reliability at a time of significant growth correlates directly with the investment and operating cost required to deliver the services. The economics of this model are not sustainable for CSPs faced with the continued growth in connections and mobile traffic created by mobile video and Internet of Things (IoT) devices.

This paradigm shift in end-user usage is driving CSPs to consider vEPC as a priority to transform networks to achieve the economic benefits of NFV.

The market opportunity for vEPC is largely based on the timing of NFV adoption and the deployment of services that support profitable business growth. A recent Infonetics report forecasts "Communications Service Provider mobile core and EPC software revenue to grow from $74M in CY14, to $218M in 2015 and over $2B in CY19 at a 5-year CAGR of 95%."[1] Figure 1 shows data from the 2015 Infonetics Carrier-SDN-NFV-HW-SW-Analysis, which provides some perspective on the software revenue projections for the mobile core.

	Worldwide Service Revenue (US$M)			2014-2019
	2014	2015(E)	2019(E)	CAGR
PCRF and DPI Functions	$553	$1,030	$3,153	42%
Mobile Core and EPC Functions	$74	$218	$2,088	95%
IMS Component Functions	$89	$396	$1,719	81%
Security Functions	$37	$79	$421	63%
vRouters	$2.5	$21	$271	156%
Video CDN Functions	$0.0	$9.9	$301	N/A
Other	$1.5	$15	$688	242%
Total VNF Revenue	**$758**	**$1,767**	**$8,642**	**63%**

Figure 1. Mobile core market opportunities.

	Worldwide Service Revenue (US$M)			2014-2019
	2014	2015(E)	2019(E)	CAGR
NFV	**$951**	**$2,264**	**$11,602**	**65%**
Hardware	**$153**	**$364**	**$1,806**	**64%**
NFVI Servers, Storage, and Switches	$153	$364	$1,806	64%
Software	**$177**	**$1,847**	**$9,409**	**65%**
NFV MANO	$13	$79	$768	125%
VNF	$758	$1,767	$8,642	63%
PCRF and DPI Functions	$553	$1,030	$3,153	42%
Mobile Core and EPC Functions	$74	$218	$2,088	95%
IMS Functions	$89	$396	$1,719	81%
Security Functions	$37	$79	$421	63%
vRouters	$2.5	$21	$271	156%
Video CDN Functions	$0	$10	$301	N/A
Other	$1.5	$15	$688	242%
Software	**$27**	**$53**	**$387**	**71%**
Outsourced Services for NFV Projects	$27	$53	$387	71%

Figure 2. EPC market revenue projections.

The figures above reinforce expectations that the Mobile Packet Core opportunity is sufficient to drive CSP investment. It should be noted that Figure 2 considers only the Mobile Core and base EPC Functions (Mobility Management Entity (MME), Serving Gateway (SGW), Packet Data Network Gateway (PGW), Serving GPRS Support Node (SGSN), and Gateway GPRS Support Node (GGSN)). It does not include other EPC, Policy and Charging Rules Function (PCRF), or Deep Packet Inspection (DPI) functions that are commonly considered part of the EPC.

With these projected revenues over the next few years, it is expected that the CSP's investments will include PCRF, DPI and additional network functions, such as Evolved Packet Data Gateway (ePDG), as part of the EPC. These network functions account for over 60 percent of the projected VNF revenue through 2019.

The nature of the challenges, such as the unique characteristics of the EPC network functions, the significant reliability requirements, and the complex integration and customization requirements of these functions with the CSP network limits the number of qualified suppliers. As a result, traditional partners will be challenged to transform both traditional network and business models. For the suppliers, it requires a transformation from a hardware model for these network functions to a software and services model.

The more common suppliers used by CSPs for EPC network functions are listed in Table 1. Please keep in mind this is not an exhaustive list, and not all of the suppliers have a production virtualized EPC offering available today.

Table 1. EPC suppliers.

Company	Traditional / Virtual	Additional information on product offerings
Affirmed Networks	vEPC only	http://www.affirmednetworks.com/products-solutions/virtualization/
Alcatel-Lucent	Traditional EPC	https://www.alcatel-lucent.com/solutions/ip-mobile-core http://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2014/10743-alcatel-lucent-virtualized-epc-delivering-the-promise-nfv.pdf
Brocade	and vEPC	http://www.brocade.com/en/products-services/mobile-networking/vepc.html
Cisco	vEPC only	http://www.cisco.com/c/en/us/solutions/service-provider/virtualized-packet-core/index.html http://newsroom.cisco.com/press-release-content?type=webcontent&articleId=1601303
Ericsson	(Connectem acquisition)	http://www.ericsson.com/ourportfolio/telecom-operators/virtual-evolved-packet-core http://www.ericsson.com/news/151029-ericsson-and-sk-telecom_244069644_c
Hitachi	Traditional EPC	http://www.hitachi-cta.com/mtc-m2m http://www.hitachi-cta.com/solutions/products
Huawei	and vEPC	http://www.huawei.com/uk/products/core-network/singleepc/index.htm http://www.huawei.com/uk/solutions/broader-smarter/hw-001548.htm
Mitel	Traditional EPC	http://www.mitel.com/evolved-packet-core
NEC	and vEPC	http://www.nec.com/en/global/solutions/tcs/vepc/ http://www.nec.com/en/global/solutions/tcs/vepc/usecase2.html http://www.nec.com/en/global/solutions/tcs/pdf/vEPC_WP.pdf
Nokia Networks	Traditional EPC	http://networks.nokia.com/fr/portfolio/products/evolved-packet-core http://networks.nokia.com/portfolio/liquidnet/liquidcore
Samsung	and vEPC	http://www.samsung.com/global/business/networks/core-network/ http://www.samsung.com/global/business/networks/insights/news/samsungs-virtualized-core-solution-chosen-to-support-sk-telecom
ZTE	Traditional EPC	http://wwwen.zte.com.cn/en/solutions/core_network/packet_core/ http://wwwen.zte.com.cn/en/press_center/news/201510/t20151022_445115.html

Industry Challenges

There are a number of challenges with the physical appliance-based EPC deployment model that are providing motivation to top CSPs to move forward with vEPC. For those CSPs the motivation includes:

- **Limited innovation**. Proprietary hardware platforms are not interchangeable from different vendors and are designed only to support that specific network function. Current EPC functions on proprietary hardware require specific operational skill sets and result in vendor lock-in, limiting innovation and supplier flexibility.
- **Capacity/Scaling**. CSPs currently dimension and build out EPC network functions for peak usage based on an 18-month forecast. This approach does not lend itself to cost or resource efficiency and results in overspending for lab resources, software licenses, and hardware capacity with substantial amount of underutilized capacity. In contrast, standard servers can be deployed at short notice to handle Operational, Marketing, or other Business requirements. The classic approach will not meet the competitive requirements of CSPs, compared to others exploiting virtualized solutions.
- **Improve end-user experience**. CSPs are evaluating the decomposition of EPC functions and considering technology that will push functions closer to the edge of the network.[2] A common Intel® Architecture-based infrastructure from edge to core enables software portability across a common architecture. Decomposition and distribution of mobile core functions enables improved performance (latency, jitter, and so on) for critical CSP revenue-generating services.
- **Service agility**. Virtualized software to enable CSPs to provide an IP-based mobile infrastructure that will scale to optimally deliver services, and allow operational flexibility for workload placement at various production or backup sites.
- **Reduce cost**. Intel® Architecture-based servers and technology deployed as the common horizontal platform for all network functions will reduce total cost of ownership (TCO) for CapEX and OpEX, through the use of standardized, reusable infrastructure instead of dedicated hardware appliances.

State of the Industry

The exact market size and growth predicted for network connected devices fluctuates; however, most experts agree it's a rapidly growing segment that requires new architectures and approaches that will scale to meet the projected traffic demands.[3] The need for change is widely acknowledged across the industry; however, the EPC is a complex and critical component of the CSP's network. Therefore, the transition planning required to implement a completely new architecture is a substantial undertaking. Also, any changes that negatively impact existing services have a direct impact on consumers and the CSP's revenue stream and customer satisfaction profile.

The mobile architecture and method by which connections are established allows for the connections and traffic to be routed to different service domains. Routing based on Access Point Naming (APN) enables CSPs to launch services in an isolated, contained domain. This provides an environment that does not impact existing services and will allow the validation of new vEPC architecture as the operational challenges and the integration with existing operations and business support systems (OSS/BSS) are worked through successfully.

The ability to provide machine-to-machine (M2M) and Internet of Things (IoT) devices with a unique APN enables CSPs to steer traffic to a specific packet core, which is isolated from those currently providing consumer and enterprise services. This approach provides an ideal opportunity for CSPs to deploy vEPC network functions without impacting existing consumer or critical enterprise offerings.

For these reasons, the majority of current vEPC deployments underway are to support IoT/M2M devices, which are growing dramatically in number and application complexity. Virtual EPC will tend to be deployed in parallel to the main production EPC and will focus on M2M and customized enterprise services.[4] Deploying vEPC for IoT provides a transition for suppliers as well. Deployment will not immediately impact their existing footprint of legacy physical network functions and will, therefore, enable continued expansion of the existing footprint into new markets, at least until Physical Network Function (PNF) products are end of life. This approach also enables new entrants and non-traditional EPC suppliers a "green field" opportunity. Eventually, the new dynamics will alter the technological and competitive dynamics of the industry.

Some examples of early vEPC deployments to support IoT that are publicly disclosed are listed here:

Affirmed Networks Selected by TELUS to Support Internet of Things Communications Service Offerings

http://www.prnewswire.com/news-releases/affirmed-networks-selected-by-telus-to-support-internet-of-things-communications-service-offerings-300186585.html

Saudi Telecom Company Goes Live with Affirmed Networks for M2M Deployment

http://www.prnewswire.com/news-releases/saudi-telecom-company-goes-live-with-affirmed-networks-for-m2m-deployment-300175855.html?tc=eml_cleartime

AT&T to Utilize Cisco Virtual Mobile Network Technology Within Its Connected Car Services

http://newsroom.cisco.com/press-release-content?type=webcontent&articleId=1601303

Common EPC and SDN Policy Initiatives

While the current benefits of virtualized EPC are well recognized across the industry, the opportunities and challenges created by innovative development cycles and complex technology integrations of vEPC with SDN capabilities loom on the horizon. To help the industry move forward, the 3G Partnership Program (3GPP) has initiated a study that examines splitting the EPC into the control and the user plan. This study delves into all the complexities, protocols, interfaces, and encapsulation requirements of EPC. Discussions as part of this study will help influence the alignment with and/or modification of the capabilities of existing SDN controllers and switches based on OpenFlow* or other approaches. Aside from the 3GPP efforts, SDN-focused groups, such as Open Network Foundation (ONF), are driving initiatives in this area.

In addition to the basic protocol alignment, integration of SDN with mobile networks will require the integration of policy and charging. Mobile networks have comprehensive policy and charging controls and a complex policy enforcement architecture (for details, refer to the 3GPP Technical Specification 23.228). The 3GPP is also conducting several technical reports and studies to evaluate approaches to better integrate the dynamic capability provided by SDN into the existing mobile network policy controls architecture and traffic steering policies. As an example, 3GPP Technical Review 23.718 defines a new interface (called the St Interface) between the Policy and Charging Rules function (PCRF) and a new Service Chain Traffic Controller function (SCTCF).

As an important contributor to those initiatives, Intel is leading standards efforts and driving ecosystem solutions based on an open reference architecture that will align the ecosystem on the integration of vEPC and SDN. The results of these efforts will likely influence how and when services are deployed to support service offerings of certain core functions of the mobile network.

Future Business Models Enabled with vEPC

The virtualization of important EPC functions will allow the CSPs to disaggregate complex mobile core functions into discrete components and distribute them to specific locations. These components largely consist of control plane and data plane elements that can be configured to form unique service chains. Leveraging SDN technologies allows for dynamic service chaining in alignment with the end-user's mobile device or application requirements.

Leveraging SDN/NFV for vEPC enables new capabilities, including network slicing, which enables the definition of discrete partitions or unique service chains for mobile applications or sets of users that have specific characteristics. Network slicing is an intelligent network mechanism to better dimension network functions to realize optimal resource utilization and improve the end-user experience.[5]

The ability to offer "slices" of the network enables CSPs to design new service offerings with different business models.[6] The network slices, which can be marketed as Network as a Service (NaaS), will provide the exact service functionality needed for the specific use cases of different industries.[7] To fully realize the resource efficiency gains and service revenue potential from innovations, such as network slicing, CSPs will require maturation of the vEPC VNFs and underlying technologies (suck as service-chaining, micro-services, containers). In addition to delivering new technologies, Intel continues to invest and participate in ecosystem standards bodies, open source projects and partner enablement programs to help move the industry forward.

SDN/NFV for vEPC Network Functions

SDN and NFV are revolutionizing the industry by driving reduced cost and increased service revenue. However, the transition to NFV requires new disparate technologies to work collaboratively, allowing for operational simplicity enabling broad adoption. The maturity and integration of these disparate technologies is captured in Intel's Network Maturity Model for CSPs.[8]

As described in the previous section, traditional evolved packet core network functions provide a variety of different services, with unique protocol and scaling requirements that are driven by the combination of control, signaling, and data plane requirements. The virtualization and optimization of data plane capabilities is an area in which Intel is directly participating. Examples include contributing to data plane optimization in many open source projects (e.g. OVS, DPDK, Open Daylight, Openstack and OPNFV). Intel's investments in both open source and partner enablement continue to grow, enabling the CSPs to accelerate the transition of control plane and data plane intensive network functions to standard high volume servers (SHVS).

Examples include:

- **vEPC control plane on SHVS**. The availability of common off-the-shelf (COTS) chipsets that can provide similar or adequate performance for some EPC control plane functions compared to existing Advanced Telecommunications Computing Architecture (ATCA)-based solutions. This opportunity provides an immediate catalyst to prioritize the transition to Intel® Architecture-based SHVS.
- **vEPC data plane performance on SHVS**. vEPC requires that computational operations must be performed by system CPUs with software-based solutions or offloaded to the NIC. As a result, the perception is that the CSP would suffer lower performances than current proprietary hardware for specific EPC functions related to traffic classification, inspection, and encryption. However, chipset evolutions, Intel® Architecture-based technology development, and enhancements to Open vSwitch* (OVS*), Intel's Data Plane Development Kit (DPDK),[9] and other acceleration capabilities are maturing sufficiently to enable virtualization of data plane-based vEPC functions. A description of those efforts and the pros and cons of different approaches to accelerate data plane performance is captured here:

https://networkbuilders.intel.com/docs/open-vswitch-enables-sdn-and-nfv-transformation-paper.pdf

Intel's Role in Addressing Market Pain Points

The market adoption of technology innovation requires the business drivers for technology to solve a problem or enable a new capability. Intel has driven the ecosystem forward to make NFV commercially viable, as reflected by a significant number of production deployments being reported.[10] As noted previously, the CSP's challenge for the Mobile Core is the investment required to scale and dimension peak workloads. IoT device proliferation and the associated call models combined with the workloads of increased video usage are now the catalysts for CSPs to change from the physical EPC appliance scaling model. The legacy approach no longer makes business or financial sense. In addition, the methodical nature in which new CSP service offerings are developed and launched does not lend itself to rapid innovation, inhibiting CSPs' ability to compete with pure Internet "over the top" or cloud-based offerings.

Intel is providing the technology and contributing to the ecosystem to enable virtualized network functions and routing applications that will scale more efficiently to optimally deliver end-to-end services. A common software-defined approach to the programmability of virtualized functions and the routing between these functions provides the ability to scale and steer traffic in a more efficient manner, and the flexibility to more efficiently dimension the network based on peak workload demands.

SDN/NFV-based vEPC requires performance and dynamic programmability to achieve the full benefits of SDN/NFV. Specifically, for vEPC, Intel Architecture-based technology and ecosystem efforts will provide the CSPs with optimal resource utilization on SHVS. Intel is driving the capabilities necessary to enable the vEPC performance, management, and programmability necessary for adoption and deployment on SHVS.

Further Technology Evolutions for vEPC

The intersection of mobile networking and cloud technologies is driving several compelling initiatives across the industry. Some vEPC initiatives that are gaining traction include decomposition, micro-services and containers.[11] These disruptive technologies may significantly influence how EPC functions are architected, deployed, and consumed.

While the technologies are early in the maturity cycle, Intel is taking an active leadership role in the ecosystem-wide innovation efforts (for example, see https://clearlinux.org). The ultimate success of these cloud technologies, and its impact on current vEPC technologies, is under evaluation with ecosystem participants and will be highlighted in future solution briefs and reference architectures.

In parallel to the performance enhancements enabling the SHVS to run vEPC functions, Intel is leading or involved with the ecosystem efforts and in standards bodies to develop and introduce to the market new architectures to take advantage of SDN/NFV for EPC. These efforts not only focus on the optimization of virtualized evolved packet core functions but also represent new architectures and concepts that further decompose and distribute these functions. The goal of these efforts is to enable more flexible and efficient use of resources and distribution of Mobile Core functions.

Some examples of the ecosystem efforts that Intel is involved in include the following:

- *Understanding the Bottlenecks in Virtualizing Cellular Core Network Functions* (http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7114735)
- *Supporting Evolved Packet Core for One Million Mobile Subscribers with Four Intel® Xeon® Processor-Based Servers* https://networkbuilders.intel.com/docs/MESH_Group_Intel_EPC_TB_FINAL.pdf

Technology Overview

The following sections describe in more detail the vEPC and technologies involved.

Traditional Evolved Packet Core

The EPC is an IP-based core network infrastructure that provides packet data services to support Long Term Evolution (LTE) Radio Access Network (RAN). The EPC consists of several integrated functions that allow for convergence of legacy licensed (2G/3G/4G) and unlicensed (Wi-Fi*) radio frequencies with wire line and other alternative networks.

The three core elements of the EPC are the Mobile Management Entity (MME), Serving Gateway (SGW), and Packet Data Network Gateway (PGW). The additional functional elements are shown in Figure 3.

The EPC enables IP-based communications and services over both wireless and wire line networks. The EPC provides a common subscriber anchor for mobility, billing, policy, and charging. Since the EPC is intended to provide a common core, subscriber policy, and mobility anchor regardless of network access, additional elements and interworking functions continue to evolve to accommodate new technologies. While this applies to the integration of small cell, Wi-Fi, and future 5G access technologies, it also applies to new networking technologies. For example, standards bodies are investigating how SDN policy is communicated and integrated with EPC policy control mechanisms, which is likely to drive new functional elements.

The EPC network functions can be largely divided into two categories: control plane and data plane. Therefore the workloads required of these functions are very different and will influence how these elements are virtualized. Table 2 identifies the primary role of the EPC Mobile Core network functions with respect to control plane versus data plane.

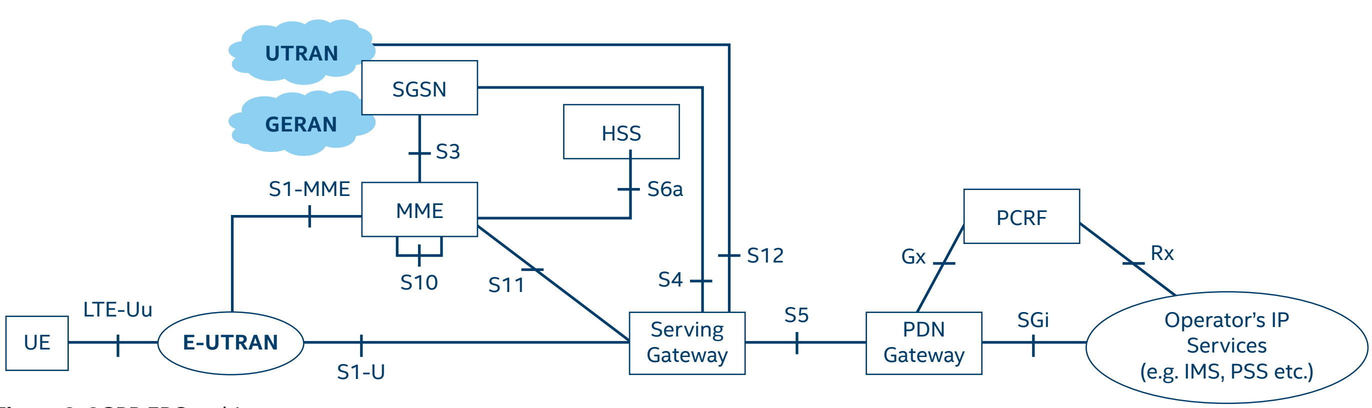

Figure 3. 3GPP EPC architecture.

Table 2. EPC control plane versus data plane functions.

Bearer/ User Plane	PGW	Packet Data Network Gateway	Note: typical PGW functionality includes DPI, policy enforcement, and other Gi-LAN network functions.
	SGW	Serving Gateway	Can be collapsed with PGW (SAE-GW)
	GGSN	Gateway GPRS Support Node	2G/3G mobile anchor
	Femto GW	Femto Cell Gateway	Small Cell (encryption requirements)
	Sec GW	Security Gateway	Can be collapsed with other functions (encryption requirements)
	ePDG	Evolved Packet Data Gateway	Wi-Fi*/EPC integration (encryption requirements)
	HA	Home Agent	CDMA or Wi-Fi mobility anchor
	WAG	Wireless Access Gateway	Wi-Fi/EPC integration
	eNodeB	evolved NodeB	Both control and user plane (S1-MME, S1-U)
Control Plane	MME	Mobility Management Entity	
	SGSN	Serving GRPS Support Node	If Fast Path used, reduced bearer plane impact
	AAA	Authentication, Authorization, and Accounting	several "types" of AAA
	PCRF	Policy and Charging Rules Function	Signaling, high TPS
	OFCS	Offline Charging System	
	OCS	Online Charging System	Signaling, high TPS
	DNS	Domain Name System	
	HSS	Home Subscriber Server	Subscription info (storage and high TPS)
	ANDSF	Access Network Discovery and Selection Function	
	SPR	Subscriber Profile Repository	Subscriber policy information (storage and transactional)
	DRA	Diameter Routing Agent	

Regardless of the control plane or data plane, EPC platform resources are used to maintain session state and signaling events. The signaling overhead for EPC functions to support events associated with mobile session set up, modification, and tear down requires significant signaling between network functions to continually exchange information. EPC functions also must maintain subscriber state, sessions, and individual flows, which has a significant impact on memory resources. The substantial transaction load and storage requirements can significantly impact system performance and scale.

The complex role of the EPC to support mobile connections requires that it scale across several unique vectors. EPC functions maintain state, signaling between functions, support substantial throughput for data plane requirements, and must scale to meet control plane requirements. This multidimensional call model and subsequent impact on the platform resources requires a unique and intelligent approach for the software to optimally use system resources (e.g., CPU, memory, I/O).

EPC Bearer Establishment

The term "multidimensional call model" is typically used to characterize the unique scaling of EPC functions and the dimensioning of those functions for capacity planning. The number of sessions, throughput, mobility events, session setup/tear down, signaling, policy, and billing interaction all impact system resources. The basic mobile session establishment requires interaction between EPC functions for authentication, policy control, billing and mobile anchoring which contribute to the multidimensional impacts on system resources.

A mobility connection typically results in multiple bearer connections for one device. These connections may have different policy, billing, and packet inspection requirements. For example, CSPs use separate APNs for basic Internet connectivity and IMS signaling for connectivity. When an IMS session is connected (for example, voice), it requires an additional dedicated bearer with specific Quality of Service (QoS) parameters driven through network-initiated, dedicated bearers. This is important because dependence on the multidimensional call model and supplier software scaling of signaling, control plane, data plane and memory will impact how the system resources are consumed as well as the ultimate capacity of the network functions and the underlying platform.

Intel Technologies and Ecosystem Enablers

For CSPs that intend to deploy vEPC solutions, Intel® Architecture and ecosystem contributions are significant. Intel's product performance, unique platform awareness capabilities, software portability from network edge to core, and contributions to open source communities and standards bodies all support solution realization.

Intel's chipset and platform capabilities enable EPC network functions to facilitate efficient resource utilization through optimal performance and programmability. Intel continues to work with the ecosystem to enable optimal use of these capabilities with seamless integration by the NFV/SDN architecture.

Virtualized network functions benefit from the ongoing efforts to enable and enhance the horizontal platform. Platform capabilities based on Intel's chipsets supporting Open Source ingredients (including DPDK[12] and OVS[13]) are leveraged by CSPs to achieve the benefits of NFV. The horizontal platform provides the foundation for a virtualized infrastructure. Capabilities such as CPU/memory virtualization, I/O virtualization, workload isolation, and acceleration are the foundation of NFV.[14]

Intel has also worked closely with ecosystem participants to develop reference architectures that maximize the value of vEPC. These architectures capitalize on open, industry-standard technologies to help CSPs reduce vendor costs; more easily produce scalable solutions; and accelerate time to market for new solutions. Purpose-built devices require CSPs and their hardware partners to qualify each version of a device, whether it is produced to offer a distinct service or to accommodate a different number of users. With vEPC based on industry-standard technologies, CSPs can produce, and qualify, fewer variations for their solutions. The virtualized environment allows them to support different services and to scale more easily.

Intel Chipset and Architecture Capabilities

Specific Intel capabilities that drive optimal performance and security for vEPC-type functions are identified in Table 3.

Some of these capabilities include Enhanced Platform Awareness (EPA)[15], Intel® Resource Director Technology (RDT), Intel® QuickAssist Technology (Intel® QAT)[16], Intel® Trusted Execution Technology (Intel® TXT), and Intel® Advanced Encryption Standards New Instructions (Intel® AES-NI) among others.

Intel® Cloud Integrity Technology features are being introduced into the processor microarchitecture to specifically address CSP requirements for secure multi-tenanted infrastructure for virtualized workloads. Specific focus areas include the security of traffic between network elements in service chains, VM-VM confidentiality, secure monitoring, ETSI compliance, and resource visibility and analytics.

These cover four layers of security:

- Platform Integrity, including Intel® TXT + TPM, AES-NI, and Secure Key
- Location & Boundary Control, including geo-tagging
- Workload Integrity
- Run Time Integrity

Table 3. Intel technologies supporting vEPC network functions

	Enhanced Platform Awareness (EPA) CPU Pinning, NUMA, Huge Pages, others	**Resource Director Technology (RDT)** CAT/CMT/MBB	**Acceleration** Codecs on Intel Architecture, Audio/Video Acceleration	**Quick Assist Technology (QAT)**	**Trusted Execution Technology (TXT)**	**Advanced Encryption Standards New Instructoin (AES-NI)**
Data Plane	✓	✓	✓	✓	✓	✓
Control Plane	✓			✓	✓	

The following table provides links to more information on these specific capabilities:

Table 4. Links to specific capabilities.

Intel® Resource Director Technology	http://www.intel.com/content/www/us/en/architecture-and-technology/resource-director-technology.html
Intel® QuickAssist Technology	http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/communications-quick-assist-paper.pdf https://01.org/packet-processing/intel®-quickassist-technology-drivers-and-patches
Intel® Trusted Execution Technology	http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/malware-reduction-general-technology.html http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-security-paper.html http://www.intel.com/content/dam/www/public/us/en/documents/guides/intel-txt-software-development-guide.pdf
Intel® Advanced Encryption Standards New Instructions	https://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/aes-ipsec-performance-linux-paper.pdf
Enhanced Platform Awareness	https://software.intel.com/sites/default/files/managed/8e/63/OpenStack_Enhanced_Platform_Awareness.pdf https://networkbuilders.intel.com/docs/openStack_Kilo_wp_v2.pdf
Open vSwitch*	https://networkbuilders.intel.com/docs/open-vswitch-enables-sdn-and-nfv-transformation-paper.pdf
Data Plane Development Kit	http://www.intel.com/content/www/us/en/intelligent-systems/intel-technology/dpdk-packet-processing-ia-overview-presentation.html https://networkbuilders.intel.com/docs/aug_17/Future_Enhancements_to_DPDK_Framework.pdf
Intel® Cloud Integrity Technology	http://www.intelserveredge.com/enhancedsecurityservers/
Hardware Offload	http://www.intel.com/content/www/us/en/ethernet-products/controllers/overview.html

Workload Placement

SDN/NFV provides the ability to programmatically control (or automate) placement of network functions and IoT services to improve user experience, provide additional security mechanisms, and scale consistent with business drivers and application requirements. The virtual infrastructure management (VIM), resource orchestration, and service orchestration provide the underlying architecture to enable VNFs to scale to meet growing workloads. Software-based network controllers interact with the underlying virtual or physical routing and switching functions to steer traffic accordingly.

The ETSI[17]-defined interface between the VIM and the infrastructure network domain (NFVi) provides the ability to discover the capabilities of the infrastructure. The orchestration environment can use this information from the VIM to determine appropriate placement of the VNF based on correlation of the VNF descriptor (VNFD) with the infrastructure capabilities. For example, the OpenStack[18] Nova libvirt driver provides a mechanism for discovering CPU instruction set capabilities and sharing with the Nova Scheduler.

Aligning the requirements of the VNF (based on the VNFD) with the infrastructure capability provides the means to achieve optimal VNF efficiency and, when combined with intelligent orchestration, provides for infrastructure efficiency. Intel refers to these attributes as Enhanced Platform Awareness.

Listed below are some Intel architecture-based capabilities that can be used by vEPC functions to optimize performance. Intel has contributed to the ecosystem to make these capabilities easily consumable and usable for workload placement decisions (for example, OpenStack). vEPC suppliers and CSPs should consider these capabilities for optimal performance and efficient data center resource utilization.

NFVI Capability Attributes Exposed to the VIM

- Non-Uniform Memory Access (NUMA) CPU and Memory configuration (collocated memory and socket).
- NUMA I/O Device Locality configuration (collocated PCI* device and socket). OpenStack will locate the virtual machine (VM) and PCI device on the same NUMA node, but only for VMs with explicitly defined NUMA topology.
- CPU pinning – OpenStack does CPU pinning if requested.
- Encryption and compression acceleration (Intel® QAT).
- Trusted Platform/Trusted Boot.
- Intel® AES-NI, Intel® Advanced Vector Extensions (Intel® AVX), SSE4.2, RD RAND (Instruction Set Extensions) – ComputeCapabilitiesFilter allows requests to land a VM on a host that has a specific CPU feature.
- vSwitches (type, capability) - OVS specified, with or without either DPDK hardware offload acceleration (DPDK/HWOA) with some DPDK awareness, not at scheduler level. Selecting vhost-user ports on supporting platforms happens behind the scenes with no HWOA awareness.
- Memory (size, page sizes, NUMA allocation policy) - OpenStack is aware of the amount of memory and page sizes present in a platform; NUMA allocation policy is supported.
- Huge Page Support (2 MB/1 GB).
- I/O Pass-through (Full PCIe* pass-through of the I/O device to the guest – Requires bootstrapping of the node in virtualized network function infrastructure (VNFI).
- I/O Pass-through (Virtual Function (SR-IOV) of the I/O device to the guest) – Requires bootstrapping of the node in VNFI). NFVI Capability Attributes not needed to be exposed to the VIM.
- CPU DDIO (Direct Data I/O).
- Network cards (interface capabilities such as LSO and LRO etc; interface BW, DPDK support).

NFVI Capability Attributes That are Under Development to Be Exposed in the VIM

- In hardware geolocation (based on Trusted Platform - due Q4 2015).
- Standard HW Acceleration API (for example, DPDK-AE).
- Last level cache utilization (CTM).
- Cache Allocation (CAT).

Open Source and Standards:

Intel is driving software contributions and broad market capabilities through important Open Source communities.

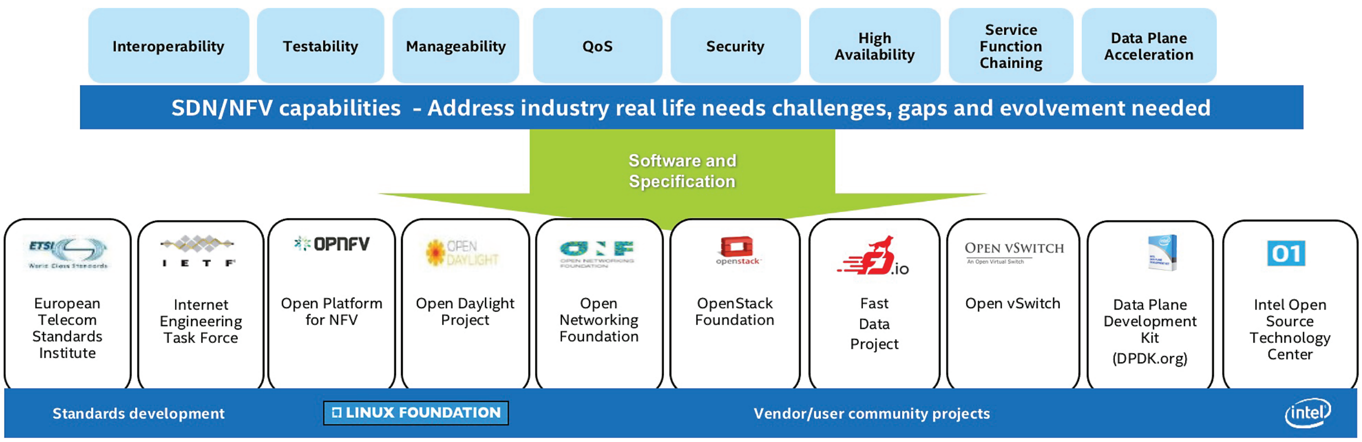

Figure 4. Intel's involvement in open source and standards.

Intel invests in 10 Open Source and Standards initiatives shown on Figure 4, from the ETSI-NFV group to Intel's own packet processing project on 01.org.

Contributions are driven both by the market and by specific customer requirements. These include commercial deployments that meet business needs, support targeted performance metrics, close development gaps, and provide the management tools needed to ensure service levels.

Intel's contribution is across the entire spectrum, including technical specifications, code development, testing, and benchmark tools and reference platforms.

Intel is involved in a number of standards efforts related to 5G and virtualization of next-generation EPC, including the decomposition of EPC. Some of the more compelling use cases involving Intel are related to the separation of control plane and bearer plan, EPC decomposition, and integration of SDN/vEPC:

- Control and User Plane Separation of EPC nodes (3GPP TR 23.714)
- Architecture Enhancements for Dedicated Core Networks (DECOR).
- Study on New Services and Markets Technology Enablers (3GPP TR 22891)

Intel® Open Network Platform Reference Architecture

Intel Open Network Platform (Intel® ONP) Server is an enablement program with a reference architecture integrating Intel's hardware and open source software ingredients for easier ecosystem adoption. One of the key objectives of Intel ONP Server is to align and optimize key Open Community software ingredients for architects and engineers targeting high-performing SDN and NFV solutions. Intel ONP provides a convenient reference platform to evaluate the latest performance contributions for OpenStack,[19] DPDK,[20] and accelerated OVS.[21]

Intel® Network Builders

Intel recognizes that part of enabling network transformation requires a strong ecosystem of partners. The Intel® Network Builders community (www.networkbuilders.intel.com) has more than 180+ partners developing SDN/NFV solutions on Intel architecture (see Figure 5). Within this community, there are more than 30 software vendors for critical SDN/NFV use cases, including vEPC. The work of the community extends to proofs of concept, reference architectures, and trials. With the help of its ecosystem partners, Intel remains committed to the development of technology solutions and capabilities that will improve the performance of virtualized network functions for CSPs.

Figure 5. Intel® Network Builders

Next Steps

- To learn more about Intel's technology for NFV, attend the courses available in the Intel Network Builders University at https://networkbuilders.intel.com/university
- To learn more about Intel Network Builder partners for vEPC and other NFV products, visit https://networkbuilders.intel.com/solutionscatalog.
- To build a testbed using the Intel ONP Reference Architecture, download the documentation at https://01.org/packet-processing/intel%C2%AE-onp.
- To get the best security in your NFV systems, specify Intel Cloud Integrity Technology in your infrastructure and VNF procurements.
- To get the highest performance from your NFV systems, specify compatibility with the Data Plane Development Kit in your Infrastructure and VNF procurements.
- To get the highest return on investment from your NFV systems, specify use of Enhanced Platform Awareness in your Orchestration, Infrastructure and VNF procurements.

Additional Information

Related efforts in Intel:

- OpenDaylight Contribution and IETF efforts on NSH
  https://tools.ietf.org/pdf/draft-ietf-sfc-nsh-04.pdf
  https://wiki.opendaylight.org/view/Project_Proposals:Service_function_chaining
- OpenStack EPA contributions:
  https://01.org/sites/default/files/page/openstack-epa_wp_fin.pdf
  https://networkbuilders.intel.com/docs/openStack_Kilo_wp_v2.pdf
- Intel Open Network Platform
  https://01.org/packet-processing/intel-onp-servers

Intel Network Builders Related Information:

- https://networkbuilders.intel.com/docs/Intel-Virtual-VOIP-Orch-RA.pdf
- https://networkbuilders.intel.com/solutionscatalog/session-border-controller-74
- https://www.brighttalk.com/webcast/12229/181563

ETSI-Defined vIMS Proofs of Concept (POCs):

POC 27: VoLTE Service based on vEPC and vIMS Architecture

POC 11: Multi-Vendor on-boarding of vIMS on a cloud management framework

[1] 2015-Infonetics-Carrier-SDN-NFV-HW-SW-Analysis

[2] http://www.lightreading.com/mobile/5g/atandt-virtualized-mobile-core-key-to-5g/d/d-id/721124

[3] http://www.fiercewireless.com/story/ericsson-backs-away-expectation-50b-connected-devices-2020-now-sees-26b/2015-06-03

[4] http://www.lightreading.com/carrier-sdn/nfv-(network-functions-virtualization)/the-rise-of-virtual-epc/a/d-id/708394

[5] http://www.ericsson.com/news/151029-ericsson-and-sk-telecom_244069644_c

[6] http://www.ericsson.com/ourportfolio/telecom-operators/accelerating-iot?nav=marketcategory002

[7] http://networks.nokia.com/.../nokia_5g_systems_of_systems_white_paper.pdf

[8] http://www.intel.com/content/www/us/en/communications/service-provider-network-maturity-paper.html

[9] http://www.intel.ie/content/www/ie/en/intelligent-systems/intel-technology/packet-processing-is-enhanced-with-software-from-intel-dpdk.html

[10] http://telecoms.com/intelligence/telecoms-com-intelligence-whitepaper-33-live-deployments-which-prove-how-nfv-has-crossed-the-chasm/

[11] http://www.rcrwireless.com/20150831/software-defined-networking-sdn/att-looks-to-containers-micro-services-in-sdn-and-nfv-push-tag2

[12] http://www.dpdk.org

[13] http://openvsitch.org

[14] http://www.intel.com/content/www/us/en/virtualization/virtualization-technology/intel-virtualization-technology.html

[15] https://software.intel.com/en-us/articles/openstack-enhanced-platform-awareness

[16] https://01.org/packet-processing/intel%C2%AE-quickassist-technology-drivers-and-patches

[17] http://www.etsi.org

[18] http://www.openstack.org

[19] http://www.openstack.org

[20] http://www.dpdk.org

[21] http://openvswitch.org

Printed in USA 0516/JL/PDF ♲ Please Recycle 334427-001US