

Evolved Packet Core (EPC) for Communications Service Providers

Table of Contents

Network Topology Overview: Evolved Packet Core	2
Basic Evolved Packet Core Architecture Components	2
EPC Integration with Policy and Charging Architecture.....	5
Next-Generation EPC Architectures Efforts	5
EPC Decomposition and Separation of User and Control Plane	5
Solution Architecture for Next-Generation Networks	5
Separation of Control Plane and User Plane for EPC	6
Network Slicing	8
ONF Solution Architecture Efforts	8
Platform Reference Architectures	9
vEPC Solution Deployment Considerations	10
Performance	10
Operational	10
How, where, what, when, and how much?	11
Next steps	11

Network Topology Overview: Evolved Packet Core

Basic Evolved Packet Core Architecture Components

The Evolved Packet Core (EPC) is an IP-based core network infrastructure that provides packet data services to support the convergence of licensed (2G/3G/4G) and unlicensed (Wi-Fi*) radio technologies. The EPC also provides the capability for the integration of wireless with wireline and other alternative networks to enable IP-based communications and services over both wireless and wireline networks. The EPC provides a common subscriber anchor for mobility, billing, policy, and charging.

The three core elements of the EPC are the Mobility Management Entity (MME), Serving Gateway (SGW), and Packet Data Network Gateway (PGW). The schematics shown in Figure 1 and Figure 2 provide a simplified view of the virtual EPC (vEPC) architecture and topology.

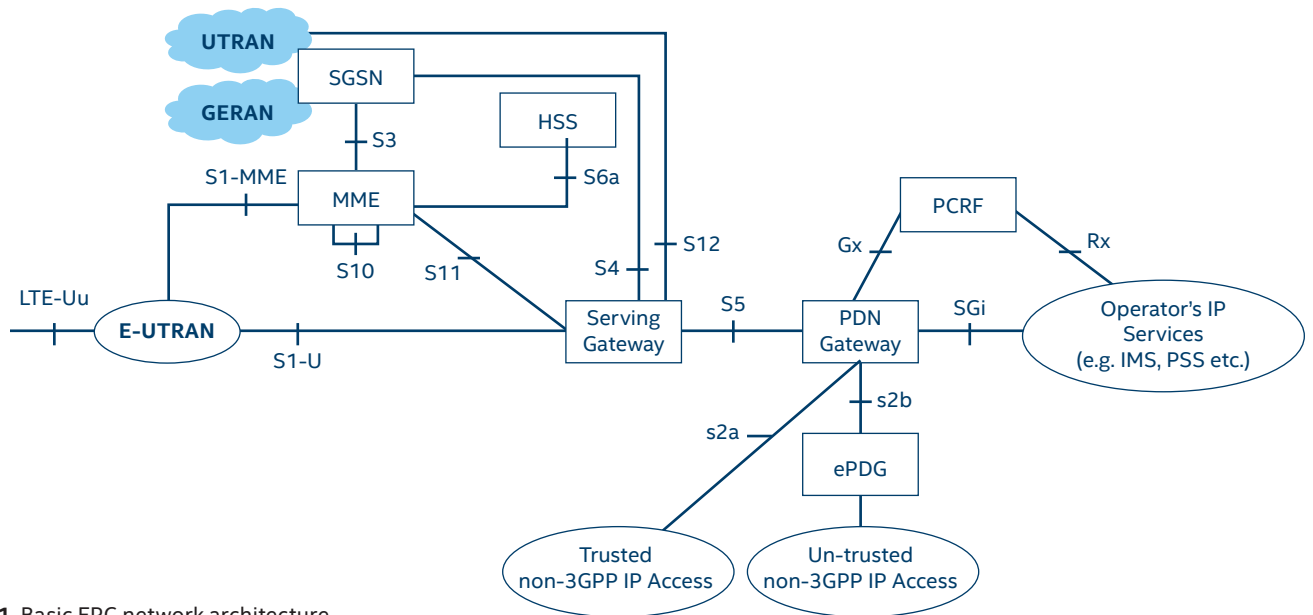


Figure 1. Basic EPC network architecture.

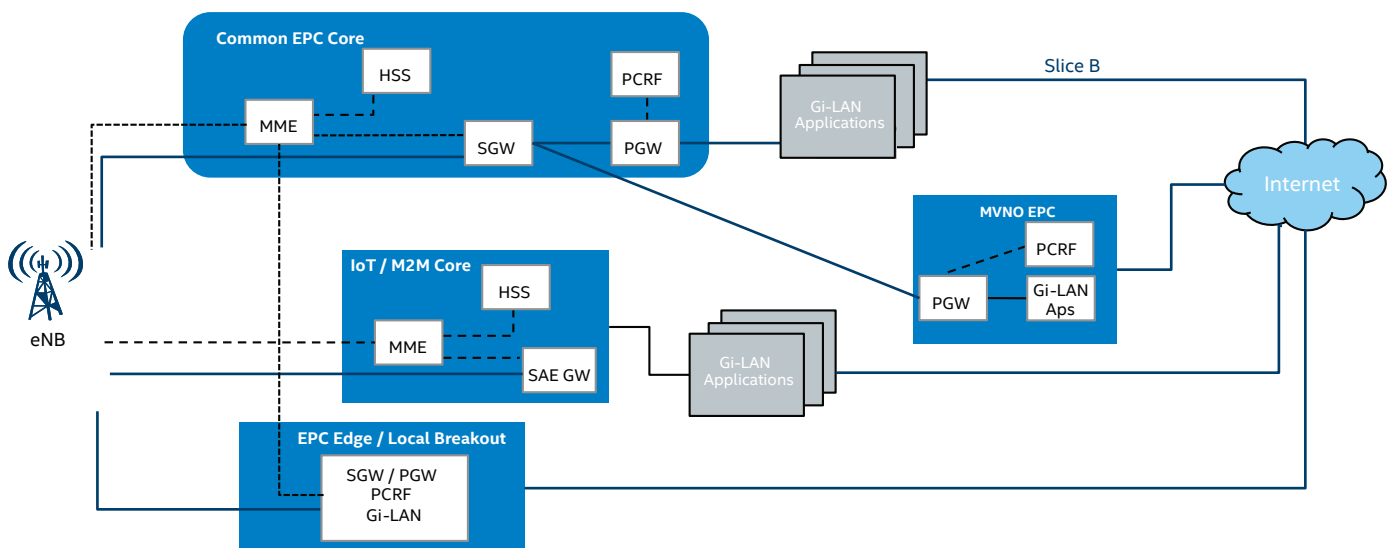


Figure 2. Basic EPC network topology examples.

The 3rd Generation Partner Project (3GPP*) defines the details of the EPC architecture, functional elements, and interface requirements. As stated above, the EPC is designed to support any on-access technology. The following links provide a good introduction to EPC-based on-access technology.

3GPP TS 23.002 (<http://www.3gpp.org/DynaReport/23002.htm>)

Provides an overview of the architecture of the 3GPP system. In particular, it describes all the network elements used in the EPC and also in legacy core networks.

3GPP TS 23.401 (<http://www.3gpp.org/DynaReport/23401.htm>)

Defines the architecture of the EPC for Evolved Terrestrial Radio Access Network (E-UTRAN) access.

3GPP TS 23.402 (<http://www.3gpp.org/ftp/Specs/html-info/23402.htm>)

Defines the architecture enhancements for non-3GPP access.

Figure 3 provides an example of the basic 3GPP interfaces for EPC to support Circuit Switch (CS) and Packet Switch (PS) services. The blue lines represent EPC-specific interfaces.

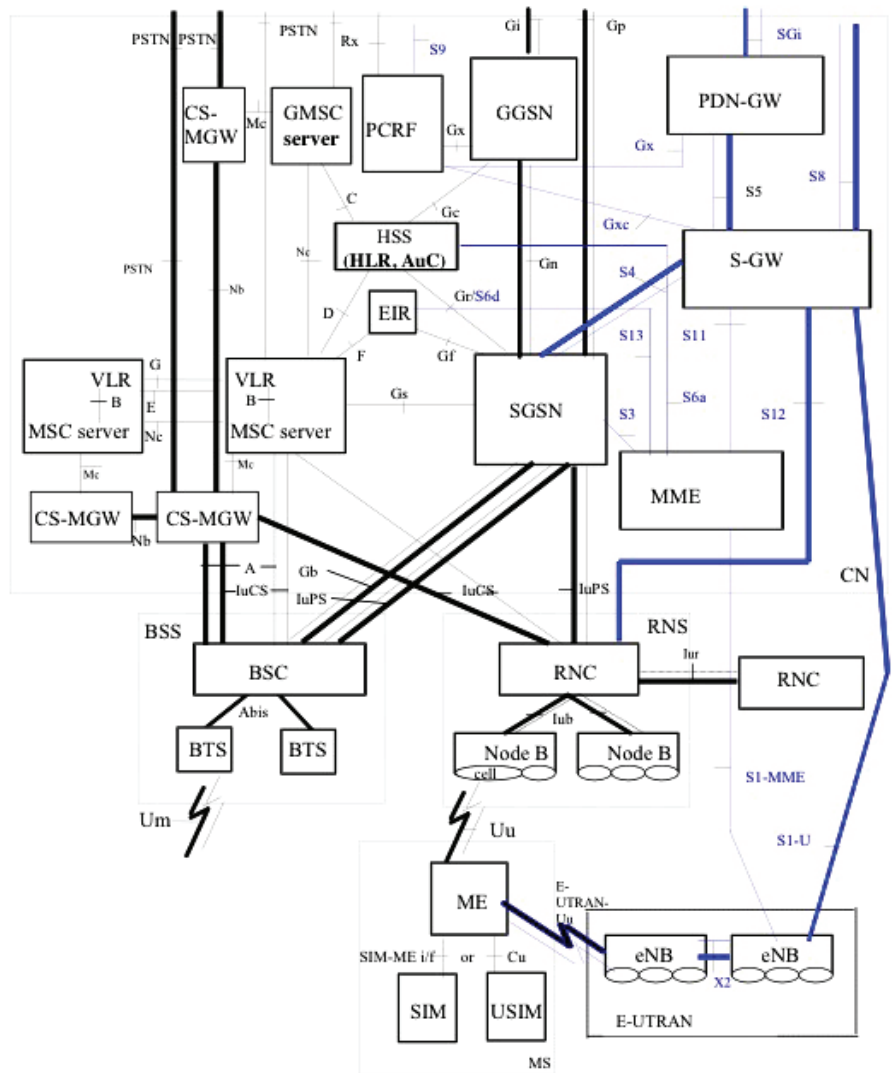


Figure 3. Basic 3GPP configuration and interfaces to support CS and PS Services.¹

Figure 4 provides an example of the basic 3GPP interfaces for EPC to support non-3GPP access. The dotted red line items represent non-macro-cellular access elements that are not in the Communication Service Providers' (CSPs) domain.

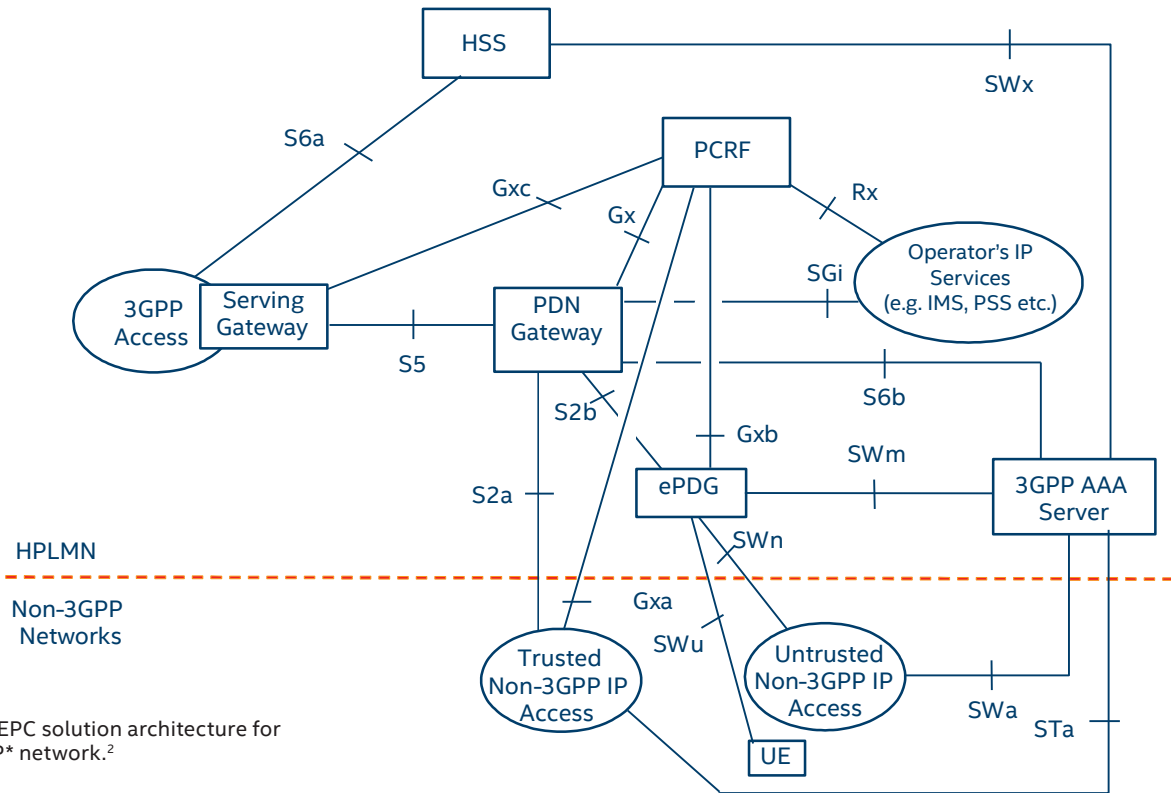


Figure 4. EPC solution architecture for non-3GPP* network.²

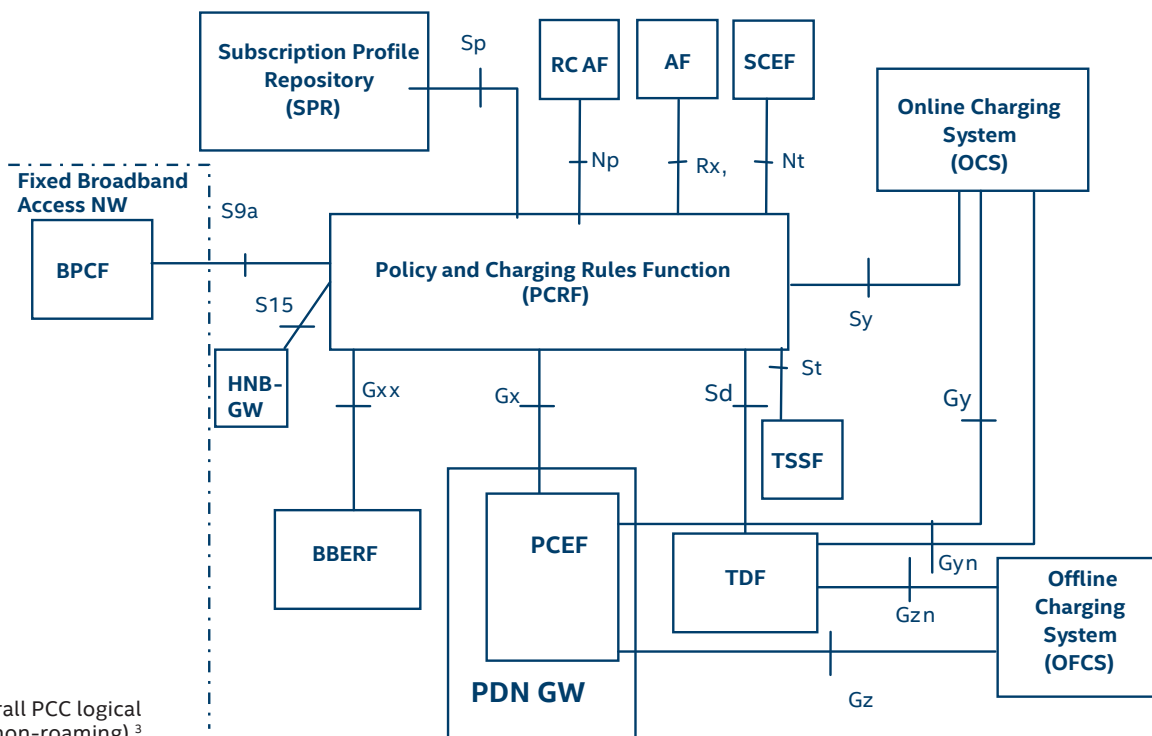


Figure 5. Overall PCC logical architecture (non-roaming).³

EPC Integration with Policy and Charging Architecture

Mobile service provider networks have comprehensive policies, charging controls and policy enforcement architecture that support 3GPP standards. These standards define the manner in which the policy functions are deployed within the CSP's network. An example of the logical policy architecture for a mobile network is shown below, as defined in 3GPP TS 23.002.

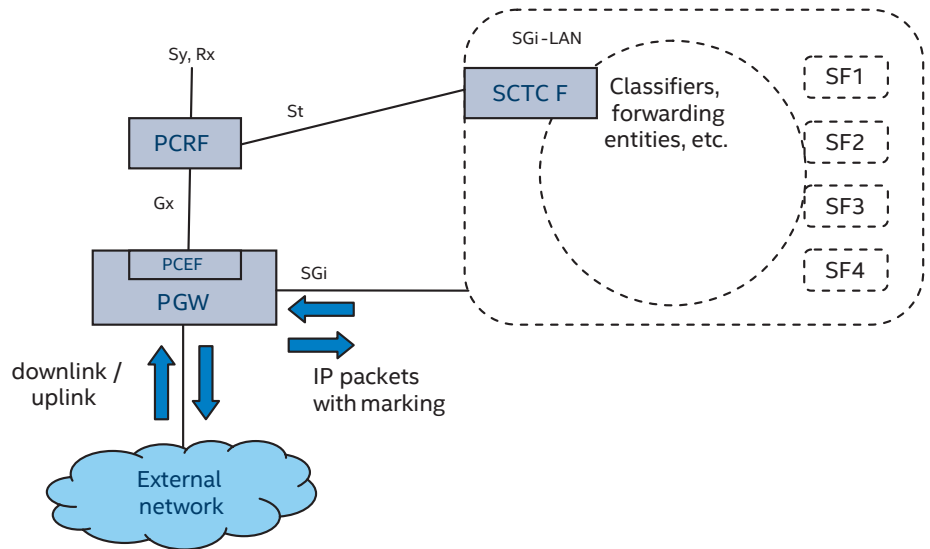


Figure 6. Example of traffic steering with packet marking in PCEF.⁴

Next-Generation EPC Architectures Efforts

3GPP-related efforts are driving solution architecture definition for next-generation software-defined network/Network Function Virtualization (SDN/NFV)-based EPC architecture. These efforts include the following:

EPC Decomposition and Separation of User and Control Plane

- 3GPP Dedicated Core Networks (DECOR)
- 3GPP Enhanced Dedicated Core Networks (eDECOR)
- 3GPP Control and User Plane Separation (CUPS)

Solution Architecture for Next-Generation Networks

- 3GPP Rel-14 study SMARTER; TR 22.891
 - 22.861 Massive Internet of Things (IoT)
 - 22.862 Critical Communications
 - 22.863 Massive Broadband

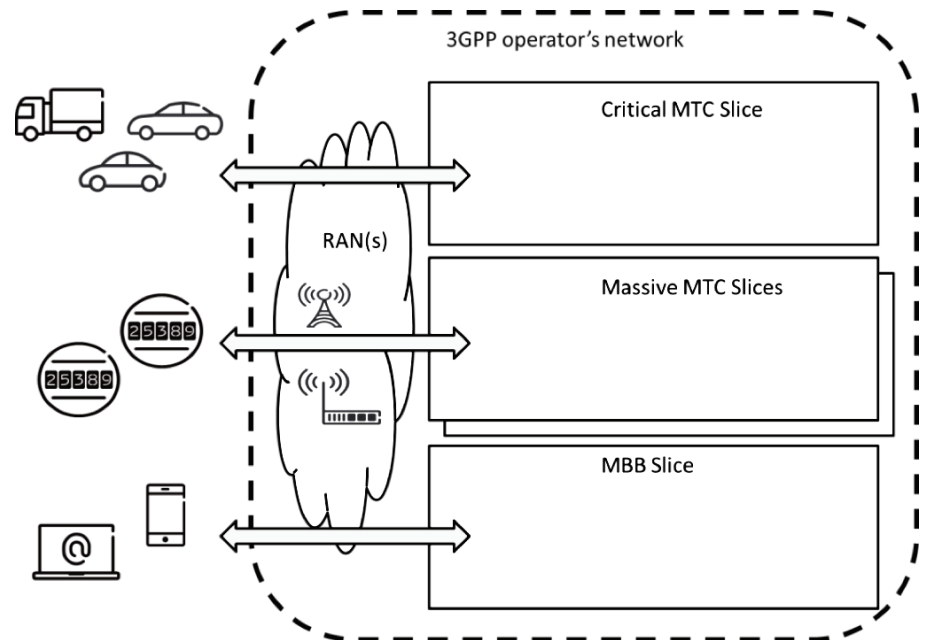


Figure 7. 3GPP network slicing for different use cases.⁵

Separation of Control Plane and User Plane for EPC

The EPC was designed to separate control plane and bearer/user (data) plane requirements to specific function elements. Some of the more common EPC nodes grouped as control plane or data plane elements are shown in Table 1.

Some EPC nodes require simultaneous, multi-dimensional scaling of both data and control/signaling planes. This is especially true for EPC data plane elements that interface with the authentication, billing, and policy and charging functions.

Figure 8 provides an example of the signaling and bearer interfaces required for the SGW, PGW, and TDF.

Leveraging the concepts of SDN, work is underway to further modify the existing vEPC architecture to enable a more comprehensive separation of control plane and data plane. This work is captured in the 3GPP technical study on CUPS.⁶

This will enable functions to scale in a more predictable manner, while providing the capability to distribute network functions more efficiently when using NFV. While this is still a technical study in the 3GPP, the industry is quickly migrating to such an architecture. This will accommodate EPC decomposition to scale more efficiently for business drivers, such as IoT and video services. The figure below provides an example of this separation.

Table 1. Common EPC nodes grouped by plane.

Bearer/ User Plane	PGW	Packet Data Network Gateway
	SGW	Serving Gateway
	GGSN	Gateway GPRS Support Node
	Femto GW	Femto Cell Gateway
	Sec GW	Security Gateway
	ePDG	Evolved Packet Data Gateway
	HA	Home Agent
	WAG	Wireless Access Gateway
	eNodeB	evolved NodeB
Control Plane	MME	Mobility Management Entity
	SGSN	Serving GRPS Support Node
	AAA	Authentication, Authorization, and Accounting
	PCRF	Policy and Charging Rules Function
	OFCS	Offline Charging System
	OCS	Online Charging System
	DNS	Domain Name System
	HSS	Home Subscriber Server
	ANDSF	Access Network Discovery and Selection Function
	SPR	Subscriber Profile Repository
DRA	Diameter Routing Agent	

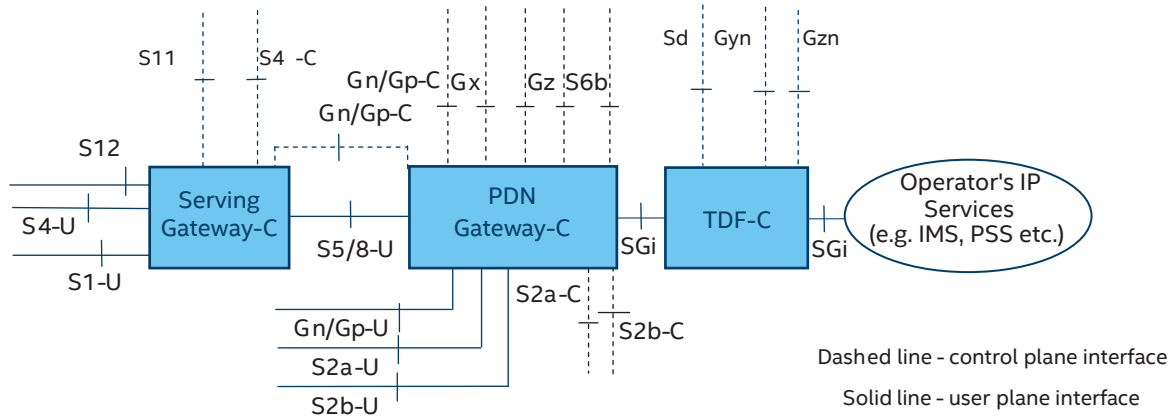


Figure 8. SGW, PGW, and TDF reference points.

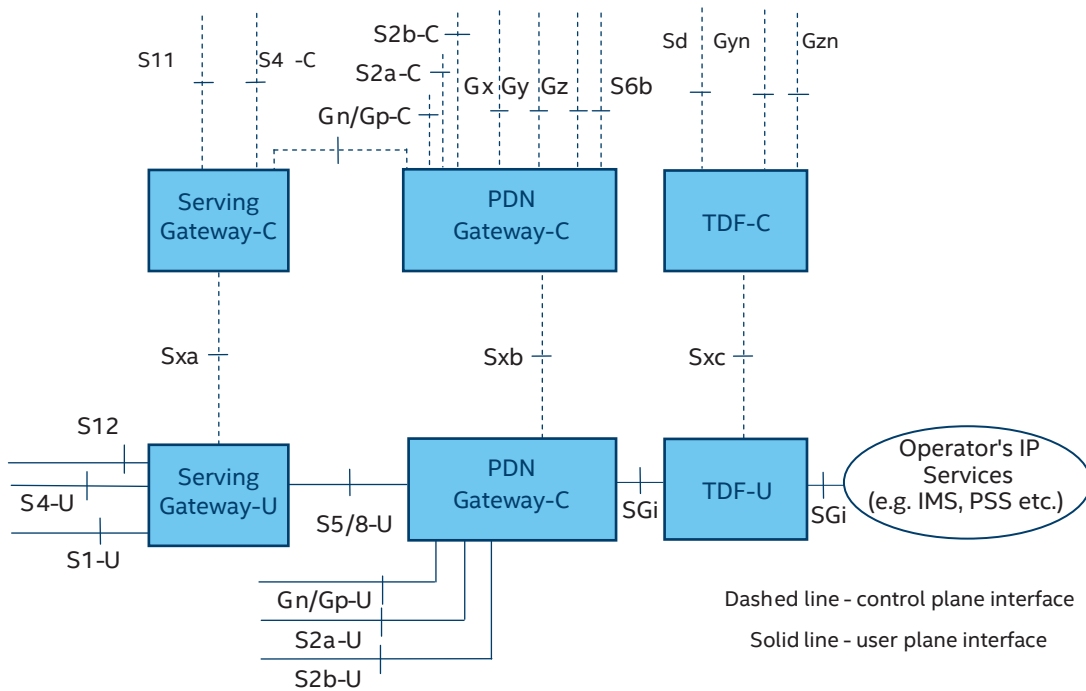


Figure 9. SGW, PGW, TF reference points (U/C separation).

Network Slicing

New revenue-generating services are being realized through the “network slicing” of the EPC. The business driver for this concept is the dramatic proliferation of IoT devices, supported by the capabilities of SDN/NFV technologies.

Network slicing is the term used to describe the partitioning (or slicing) of the network functions into more efficient segments to support a greater diversity of use cases. Figure 10, from the Next Generation Mobile Network 5th Generation (NGMN 5G) white paper, illustrates the network slicing concept.

Applying service function chaining (SFC) with SDN/NFV capabilities optimizes the EPC’s ability to steer traffic to specific network slices. SFC is described in Intel’s Gi-LAN/ SFC Solution Reference Architecture document.

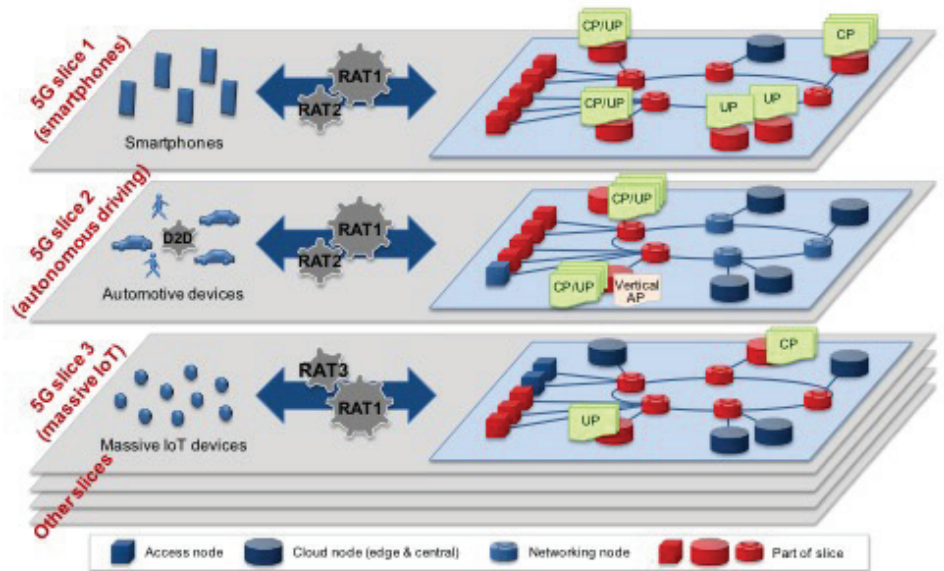


Figure 10. NGMN network slicing illustration.

ONF Solution Architecture Efforts

The Open Networking Foundation (ONF) has established a mobile core project focused on defining the EPC control and data plane requirements and architecture for an OpenFlow*-based implementation. The effort is being used to define any new changes to the OpenFlow protocol that would be necessary to support an implementation of EPC.

Figure 11 shows an SDN-based EPC with control plane and data plane separated with OpenFlow extensions.

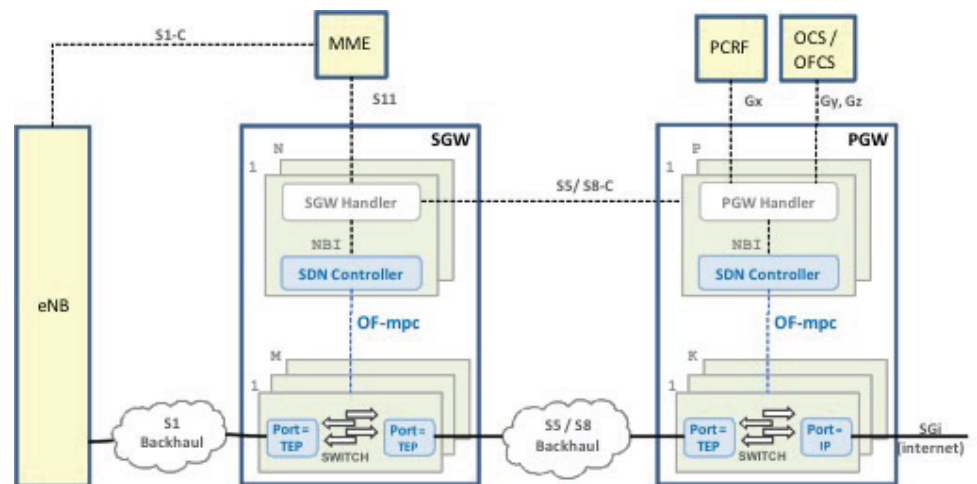


Figure 11. NGMN network slicing illustration.

Platform Reference Architectures

Table 2 shows an example of a hardware bill of materials for a vEPC reference architecture.

Additional technologies, including Enhanced Platform Awareness (EPA),¹⁰ Intel® Resource Director Technology,¹¹ Intel® QuickAssist Technology (Intel® QAT),¹² Intel® Trusted Execution Technology (Intel® TXT),¹³ and Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI),¹⁴ among others, are all relevant to the vEPC reference architecture. Please note that not all processors support all relevant technologies.

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

Table 2. Example reference architecture hardware bill of materials.

Processor	Dual Intel® Xeon® processor E5-2680 v3	
Platform	Intel® Server Board S2600WT 3 TB HDD Seagate Constellation* CS 2x Intel® Solid State Drive Data Center S3700 Series/200 GB	
Memory	256 GB Memory	
NICs	2x Intel® Ethernet Converged Network Adaptor X520-DA2	Niantic 2-port 10Gb. For the latest on controllers see Intel® Ethernet Controller
Add-In	Intel® QuickAssist Adapter 8950	PCIe* Gen 3 Refer to Drivers for Intel® QuickAssist Technology for the latest drivers and patches

Table 3. 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
Hardware Offload	http://www.intel.com/content/www/us/en/ethernet-products/controllers/overview.html

vEPC Solution Deployment Considerations

Performance

For all virtual applications, there is a need to have adequate performance on Intel® architecture-based standard high volume servers (SHVS) to justify the fundamental business case when transitioning from physical to virtual-based network functions. The ability to run more virtualized applications on the same physical hardware enhances the business case for SHVS.

However, the wide-ranging functionality of the disparate vEPC Virtualized Network Functions (VNF) requires the SHVS to deliver the capabilities to enable VNFs to scale across multiple application scenarios. Some examples of vEPC application requirements include Mobility Events, Signaling Terminal Proxy Server (TPS), Control Plane TPS, Data Plane Throughput, Subscriber State, and Connections per Second.

In addition to the latest Intel architecture chipset, the application environment should contain the capabilities identified in the previous section (pQoS, Intel QAT, Intel AES-NI, Intel TXT, EPA, Open vSwitch* (OVS), Data Plane Development Kit (DPDK), Media Acceleration, and so on) to assure optimal data center resource utilization. This applies not only to the capabilities available in the infrastructure but also to the specific capabilities envisioned for the virtual applications.

Lab trials and proofs-of-concept (PoCs) continue to demonstrate and allow for the evaluation of different paths to reach optimal performance for virtualized functions on SHVS. Standard kernel-based OVS is not sufficient to meet the performance requirements to prove the business case for SHVS. While single root I/O virtualization (SR-IOV) enables better performance than standard OVS, it has certain dependencies (for example, hardware and specific drivers) that create an operational challenge. In addition, its architecture does not ultimately deliver the benefits of virtualization. However, results from recent PoCs and lab trials of DPDK-enabled OVS (user space OVS) have delivered performance results very near that of SR-IOV.

With this in mind, the recommended reference architecture to support the business case for the SHVS replacement of dedicated physical appliances must deliver an environment that supports SR-IOV and DPDK-enabled OVS. This approach

enables CSPs to get to market immediately on SHVS using DPDK-enabled OVS or SR-IOV, if required. The transition to DPDK-enabled OVS will support the realization of the benefits of virtualization, such as optimization and specific feature sets.

Operational

As identified in the previous section, the successful implementation of a virtualized vEPC solution requires coordination of policy across 3GPP and SDN domains. However, as part of the decision process for the policy framework, it also requires the virtual network services to provide integrated hooks across all layers of the stack. As shown in the figure below, horizontal solutions must link capabilities across all the disparate layers of the NFV/SDN stack to actually deploy, monitor, and bill for services. To optimize the benefits contributed by the vEPC, there must be tight coordination of the technology environment and resource models with the CSP's subscription plans and policy entitlements

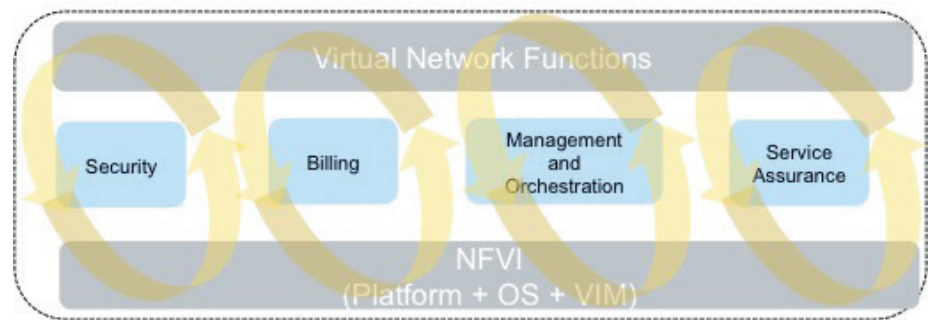


Figure 12. Horizontal solution linkage

How, where, what, when, and how much?

There are important questions that must be addressed to effectively design, integrate and operate a virtualized vEPC (for example, how, where, what, when and how much). The ability to deploy VNFs successfully on SHVS requires both an understanding of the environment's capabilities and location and tight coordination with the application's requirements. SHVS provides significant flexibility and workload placement from the edge to the core of the network. SHVS also provides the operational tools that will ensure the environment is trusted, secure, and has the appropriate resources for the task.

As an example, Intel is driving industry initiatives and making contributions, such as OpenStack Enhanced Platform Awareness (EPA), to ensure that Intel architecture-based platform capabilities are exposed sufficiently to the CSP's management and orchestration implementations to enable optimal data center resource utilization.

Network service life cycle management, monitoring, security, and billing models are all examples of the horizontal solution capabilities required for success and scaling in a virtualized environment. Intel is working closely with CSPs to enable and, in some instances, drive industry efforts to accelerate maturation of these horizontal capabilities. Future versions of this document will provide more detailed examples of those efforts.

As described in Intel's Service Provider Network Maturity Model, the initial deployments of network virtualized functions will rely on single-vendor implementations for end-to-end vertical solutions that include the necessary horizontal capabilities. As the horizontal capabilities mature and service providers transition from single-vendor implementations, the business case for creating commercial business services from the virtualization of network functions will continue to improve and will also help drive new revenue opportunities.

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 vGiLAN 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.



¹ <http://www.3gpp.org/DynaReport/23002.htm>

² <http://www.3gpp.org/DynaReport/23402.htm>

³ <http://www.3gpp.org/DynaReport/23002.htm>

⁴ <http://www.3gpp.org/DynaReport/23718.htm>

⁵ <http://www.3gpp.org/DynaReport/22891.htm>

⁶ <http://www.3gpp.org/DynaReport/23714.htm>

⁷ https://www.ngmn.org/uploads/media/NGMN_5G_White_Paper_V1_0.pdf

⁸ <http://www.opennetworking.org>

⁹ <https://www.opennetworking.org/images/stories/downloads/working-groups/charter-wireless-mobile.pdf>

¹⁰ <https://software.intel.com/en-us/articles/openstack-enhanced-platform-awareness>

¹¹ <http://www.intel.com/content/www/us/en/architecture-and-technology/resource-director-technology.html>

¹² <https://01.org/packet-processing/intel%C2%AE-quickassist-technology-drivers-and-patches>

¹³ <http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/malware-reduction-general-technology.html>

¹⁴ <https://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni>

¹⁵ www.openstack.org

¹⁶ <http://www.intel.com/content/www/us/en/communications/service-provider-network-maturity-paper.html>

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

Copyright © 2016 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Printed in USA

0516/JL/PDF

Please Recycle

334471-001US