White Paper TE-524262

Empowered by Innovation

NEC Virtualized Evolved Packet Core – vEPC

Design Concepts and Benefits

INDEX

Leading the transformation into Mobile Packet Core Virtualization

Massive growth in traffic fueled by the proliferation of smart devices, the popularity of online videos and the emergence of Internet of Things (IoT) is putting an ever-increasing pressure on mobile operators' packet core infrastructure. To that end, operators are realizing the inherent challenges in delivering the economy of scale using the current packet cores built with dedicated hardware based appliances. Reducing the total cost of operation while tackling the high traffic demand is becoming critical for planning the next evolution of mobile packet cores.

Faced with the challenges of flattening Average Revenue Per User (ARPU) from being dumb pipe to Over-The-Top providers, operators are also now getting excited about new revenue generating opportunities made possible via service driven network infrastructure. An example of such services is to offer packet core infrastructure-as-a-service bundled with other value-added network services (routing, firewall, DPI, NAT etc) in a multi-tenant fashion to enterprises. Such a packet core infrastructure-as-a-service needs to provide flexible, rapid and on-demand infrastructure deployment and capacity provisioning. Proprietary, hardware appliance-based solutions cannot deliver resource virtualization and soft-infrastructure provisioning – a key requirement in any service-driven network abstraction.

Driven by the above requirements, while leveraging the maturity in the infrastructure (compute, storage, network) virtualization technology and mobile networking solutions, NEC is one of the first to introduce the commercially available virtualized Evolved Packet Core (vEPC). NEC vEPC marks a key step towards operators' much needed transformation into the network function virtualization (NFV) space as being standardized by ETSI. NEC vEPC supports the 3GPP release 11 standards, and interoperates with legacy packet core equipments.

NEC vEPC can be deployed on any virtualized off-the-shelf Intel x86 based commodity server. Decoupling the EPC system from the underlying hardware infrastructure provides the benefit of a radical reduction in infrastructure cost. The ability to leverage virtualized infrastructure avoids piling of hardware silos and enables amortization of the common resource pool through sharing across network and IT domains, thereby resulting in significantly reduced operational costs.

Beyond deployment openness and flexibility, the fundamental advantage of NEC vEPC is its ability to soft-provision and to scale elastically. NEC vEPC system components can be provisioned on-demand and can be scaled based on diverse and time-varying control and data plane traffic requirements. Furthermore, addressing the challenges associated with virtualized commodity hardware, NEC vEPC incorporates technology innovations to deliver carrier-grade reliability and performance.

NEC vEPC simplifies the operation and management of packet cores via smart automation and orchestration of the provisioning and operation procedures. The system exposes service level APIs that enable integration with external applications building on the premise of mobile packet core infrastructure-as-a-service.

vEPC System Architecture Overview

NEC vEPC provides two core network functions: vMME (Mobility Management Entity) and vS/P-GW (Serving and PDN gateway). vS/P-GW can be deployed as totally separated or co-located. These virtual network functions (VNFs) are decomposed into elementary virtual machines. The figure below shows the virtualized environment leveraging NEC's performance optimized, open, KVM-based carrier-grade hypervisor.

The vEPC is a fully integrated platform aligned with the ETSI NFV recommendation as shown below. The core part of this platform is the NFV Management and Operation (NFV M&O) system which contains VNF manager and VI (Virtual Infrastructure) manager. The VNF manager orchestrates and manages the VNF provisioning, operation automation, and lifecycle management. The VI Manager manages the underlying virtual infrastructure composed of servers, storage and network. At the top, the E2E service orchestrator provides a service-driven interface to operators for on-demand service deployment and configuration.

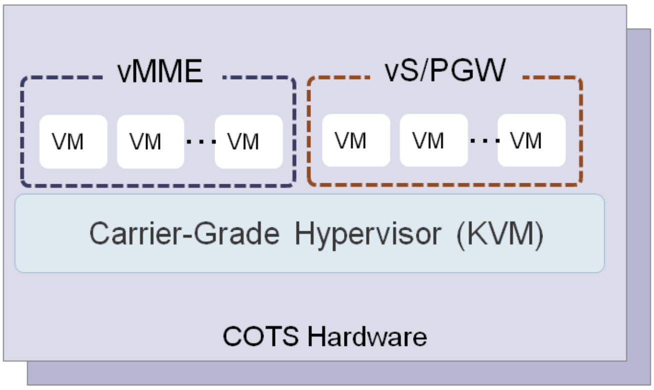

NEC vEPC Virtualized Environment

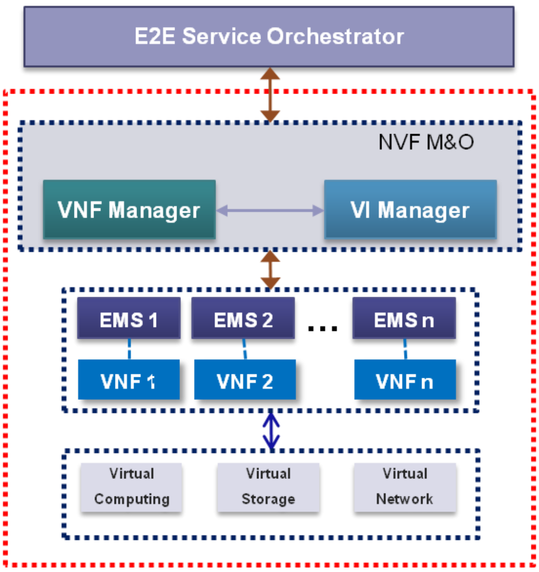

NFV architecture

The NEC vEPC is a fully integrated platform aligned with the ETSI NFV recommendations

Elastic Scalability

The unique design principle underlying the NEC vEPC system is an intelligent function decomposition that enables agility and elasticity in system scaling. At the management and provisioning level, vEPC defines logical system units which can be independently managed, deployed and scaled. At the functionality level, the decomposition separates the control plane functionalities from the data plane functionalities resulting in independent scaling of control plane and data plane. The figure below illustrates the decomposition approach.

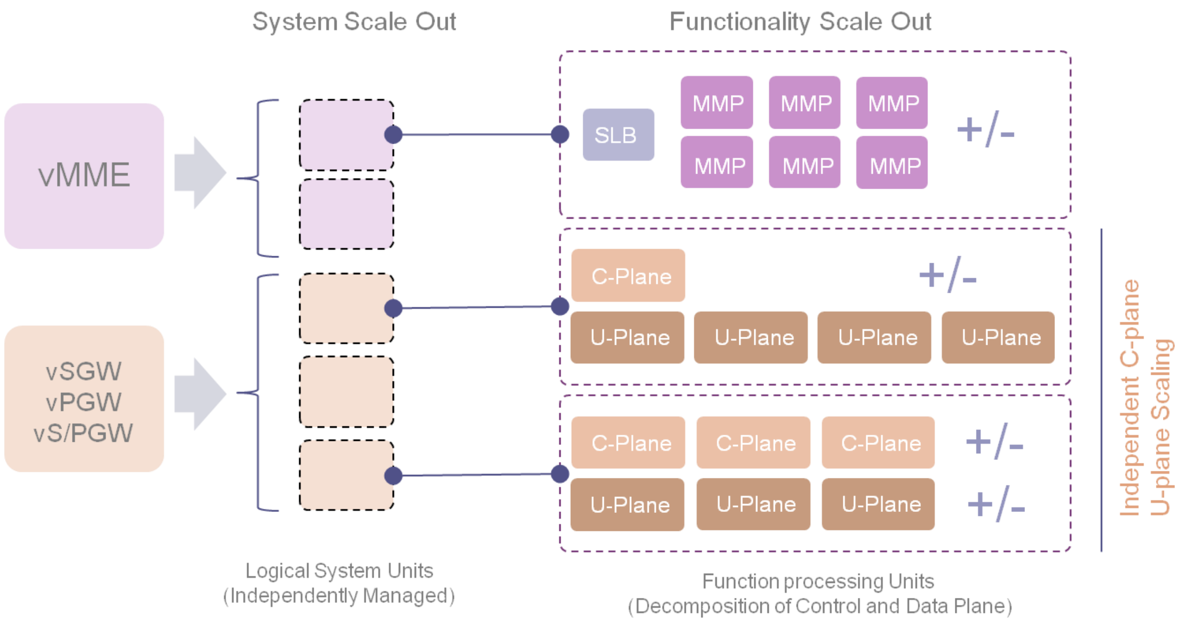

NEC vEPC features an intelligent function decomposition that enables agility and elastic system scaling

As shown in the figure, the logical system unit definitions allow scale with the management complexity and enable independence in terms of configuration to meet application requirements. With functional decomposition, each processing unit is deployed as virtual machines. In the case of vMME, the basic functional processing unit is MMP (Mobile Management Processor). vMME can be scaled by using signaling load balancers along with adding/deleting MMPs. vS/P-GW is decomposed into two main processing units: C-Plane for processing the control traffic and U-Plane for processing the data plane traffic. Each logical system unit can be independently scaled. For example, an M2M application having high C-plane traffic can have a higher number of C-Plane processing units, while a video streaming application can have higher number of U-Plane processing units. The ability to independently scale out based on traffic requirements is the core benefit of the NEC vEPC system.

NEC vEPC provides reactive (auto-scaling) and predictive (planned) horizontal scaling of resources by automating the functional processing unit provisioning through virtual machine instantiation. In case of auto-scaling, the system can scale each functional processing unit while monitoring real-time load parameters, thereby adapting to time-varying application specific traffic. Predictive or planned scaling allows operators to provision capacity bounds on-demand whereby VNF manager provisions the appropriate amount of resources to meet the capacity demand.

Availability: Redundancy and Failover

Migrating to a virtualized environment brings the challenges in terms of stability and availability of the overall system. NEC vEPC tackles these challenges through well-designed redundancy, failover and VM recovery schemes in order to meet the carrier-grade requirements. The vEPC system implements different redundancy mechanisms specific to functional processing units based on criticality while maximizing the resource efficiency. As shown in the figure below, the three different redundancy mechanisms each have a different impact on the resource requirements. In the case of the Active-Standby configuration (1+1, N+1), in-memory state synchronization enables fast failover with consistent state migration.

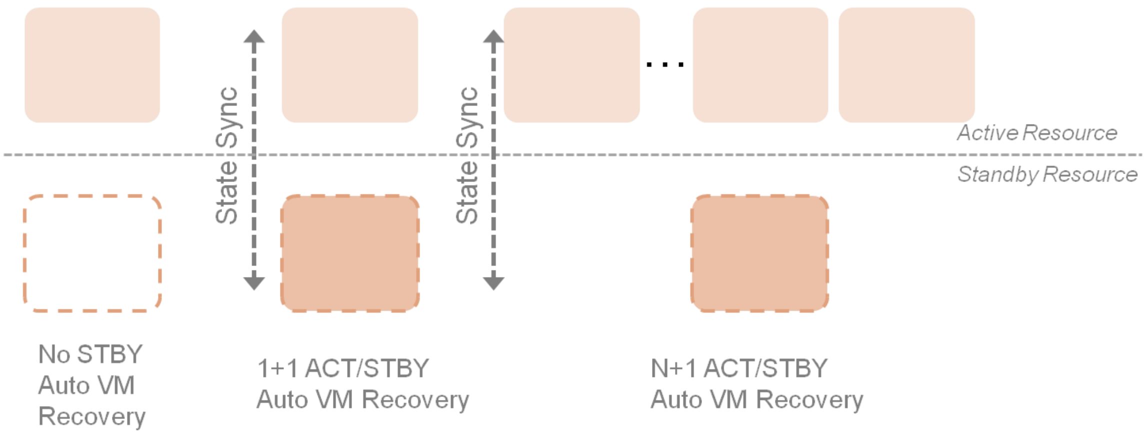

NEC vEPC's redundancy, failover and VM recovery schemes meet carrier-grade requirements

The entire process of functional failover is managed and orchestrated by the VNF manager as an auto-healing process. Firstly, NEC vEPC implements extensive virtual and physical system monitoring, logging and alert generation mechanisms at the hypervisor layer. Secondly, based on these alerts, the entire failover process is automated from the failover to the recovery of the virtual machines to replace the faulty physical or virtual machines. Furthermore, an intelligent physical resource allocator in vEPC ensures the active and standby virtual machines are located on separate physical machines to ensure high availability in case of failure. Apart from functional redundancy, NEC vEPC also adds physical and virtual network level redundancy by having dual paths for inter virtual machi ne communication.

Performance: Predictability, Stability and Optimization

Simple deployment of vEPC VNFs as VMs on a general purpose IT-grade hypervisor may not meet the high performance objectives that are required in a carrier-grade environment. Fundamental problems in an IT-grade virtualized environment are: a) Processing delay and high variability in performance due to hardware resource contention between VMs; b) Degradation of performance due to the virtualization layer overhead; and c) I/O throughput limitation by the introduction of vSwitch for network virtualization. Mobi le operators expect a predictable environment that provides certain bound on the packet processing delay as well as high I/O throughput for U-plane traffic. For example, a VoLTE call will require certain packet level delay and jitter bound from P-GWs in order to meet the end-to-end QoS requirements.

In order to address these challenges, NEC vEPC provides the Carrier Grade Hypervisor that extends the open KVM hypervisor functionality by adding priority controller for CPU, disk, and network I/O resources that are shared among VMs and by improving memory access latency; as shown in figure below. Furthermore, in collaboration with Intel, NEC vEPC leverages DPDK to provide I/O performance optimization. Additionally, NEC vEPC adds an extensive set of diagnostic and tracing tools to enhance fault management.

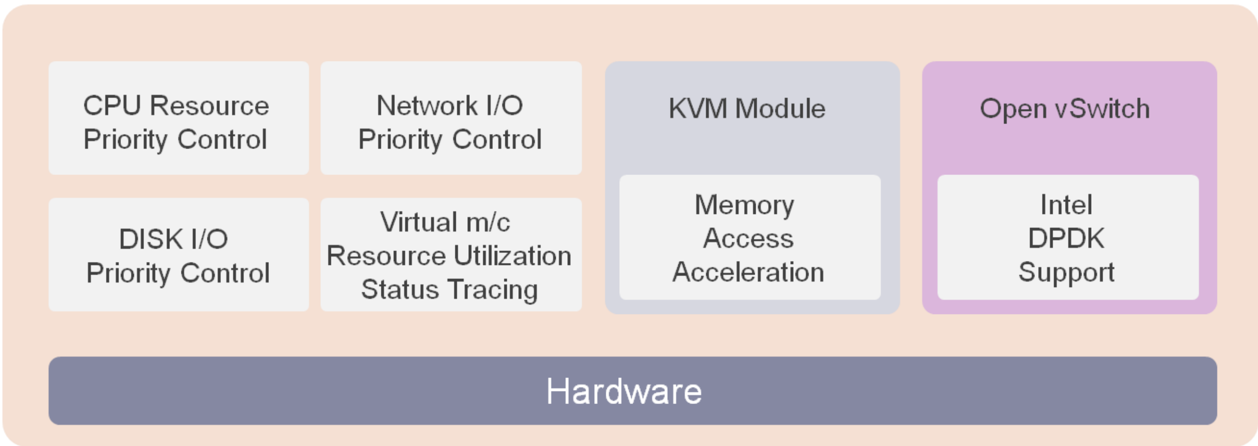

Carrier-Grade Hypervisor

A key requirement while deploying VNFs is to have a predictable virtual environment from a performance standpoint, while ensuring optimal resource usage. NEC vEPC provides optimal resource allocation through intelligent VNF placement onto physical servers. The allocation is based on the characterization of each VNF's VM resource requirements, inter-VM communication requirements and the available physical common resource pool.

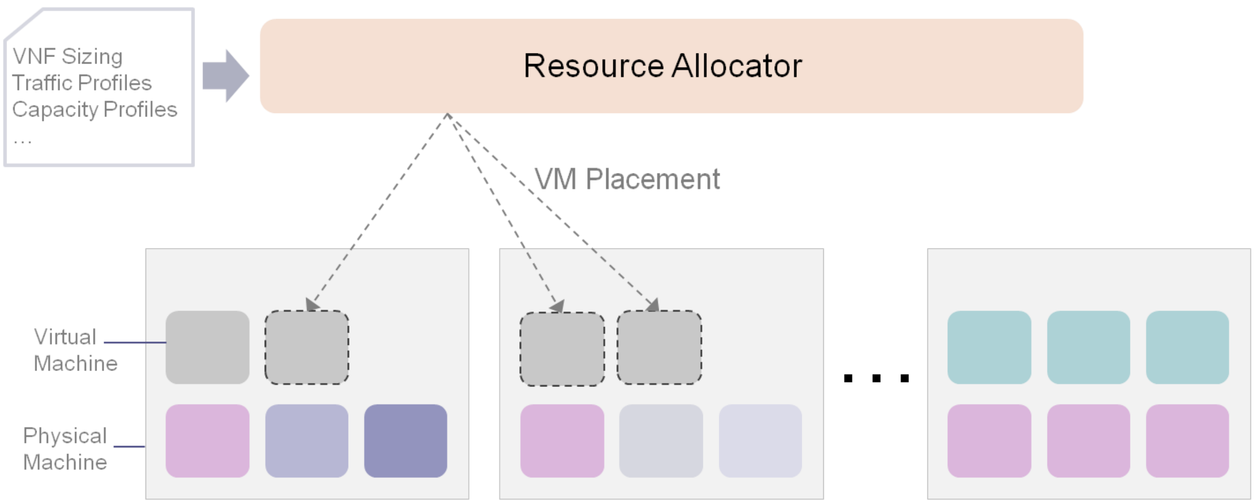

Predictable performance with high resource utilization

Orchestration: System and Infrastructure Automation and Management

The NEC vEPC VNF manager working in conjunction with the VI (virtual infrastructure) manager automates the end-to-end operation of vEPC. Key features of the automated orchestration system which minimizes operator intervention include:

Auto-scaling: Automate the scaling of resources based on dynamic capacity requirement. An example scenario is where CPU usage more than 60% for over 5 minutes, which leads to adding a new VM for the functional processing unit.

Auto-healing: Automate the entire process of system recovery based on monitoring the system diagnostics and triggering the failover plans based on the redundancy styles as well as recovering the original system to meet the workload.

Auto-provisioning: Automate the process of provisioning of the logical system units and deployment of VMs. Provisioning also includes network automation.

Auto-resource allocation: Automate the process of physical resource allocation based on performance criteria, redundancy criteria, and resource availability criteria.

The orchestration system leverages service templates that define the automation workflow for various processes. Such a system enables easy deployment to support an agile and dynamic environment and meets heterogeneous and time-varying application requirements.

Use Cases

NEC vEPC enables unique use cases that are not possible with current appliance-based EPC. One such use case is multi-tenant on-demand P-GW-as-a-service similar to infrastructure-as-a-service available from public IT cloud operators. Telecom operators can now offer their enterprise customers a service where enterprise customers can on-demand deploy P-GWs from a catalogue of PGWs with different capacity and different scalability configurations. Operators can also provision P-GWs meeting application-specific requirements, wherein P-GW with high C-plane capacity can be provisioned for M2M application while P-GW with high U-plane capacity can be provisioned for video streaming application. Furthermore, NEC vEPC can provide real-time auto-scaling to meet flash-crowd use cases such as sports video event streaming.

NEC vEPC also delivers flexibility in terms of deployment options. Operators can now deploy S-GW, P-GW or S/P-GW either at the edge (RAN) or in the Core Network (CN) or at any data center site. Deploying at RAN or CN can result in better network utilization. Deploying common DC infrastructure can result in higher amortization of resources used for multiple NFV systems. Depending on traffic profiles and network architectures, operators now have the flexibility to choose a combination of the above deployment models in reducing the overall operational cost.

Another important use case for vEPC is to deploy a disaster-resilient mobile core where resources can be pooled from geographically diverse sites to meet the urgent resource demands during a disaster scenario. Resources can also be dynamically re-allocated based on priority of different services (voice vs data) in case of such disaster.

Central to NFV is the ability to compose different VNFs to provide rich and flexible end-to-end services. With the adoption of NFV, multiple virtualized SGi functions can be dynamically added to an end-to-end path through service chaining. NEC vEPC has the unique ability to work with software defined networking (SDN) enabled networks that implement service chaining so that SGi services such as virtualized router, DPI, NAT, Media Optimizer, or Cache could be inserted selectively in the end-to-end path.

Conclusions

NEC vEPC is poised to lead the transformation of operators' mobile core infrastructure from a proprietary hardware appliance-based model into an open, virtualized and software-defined infrastructure-as-a-service model. The NEC vEPC design focuses on bringing key innovations to meet the challenges in adapting to a virtualized mobile core: availability, stability, performance, agility in service provisioning and bundling, and operational simplicity. Beyond meeting the above challenges, NEC vEPC meets the central goal of any virtualized system: radical reduction in total cost of operation through elastic scale out and efficient and dynamic resource allocation.

Abbreviations

3GPP	Third Generation Partnership Partnership Project
API	Application Programming Interface
COTS	Commercial off-the-shelf
C-Plane	Control Plane
CPU	Central Processing Unit
DC	Data Center
DPDK	Data Plane Development Kit
DPI	Deep Packet Inspection
E2E	End-to-end
EMS	Element Management System
ETSI	European Telecommunications Standards Institute
I/O	Input/Output
KVM	Kernel-based Virtual Machine
M2M	Machine to machine
NAT	Network Address Translation
P-GW	PDN Gateway
QoS	Quality of Service
RAN	Radio Access Network
S-GW	Serving Gateway
U-Plane	User Plane
VM	Virtual Machine
VoLTE	Voice over LTE