Virtualisation of Mobile Core Network Use Case
draft-king-vnfpool-mobile-use-case-02

Abstract

   Accessing the Internet via mobile data services using smartphones,
   tablets, and mobile data USB dongles has increased rapidly, as
   high-speed packet data networks provide the bandwidth required for
   today's Internet applications. Mobile operators will continue to
   evolve their core networks to the Long Term Evolution (LTE)
   Evolved Packet Core (EPC) to meet the mobility, latency and
   bandwidth requirements for mobile data users.

   Network Functions Virtualization (NFV) looks to reduce mobile core
   network complexity and related operational issues by leveraging
   standard IT virtualization technologies and consolidate different
   types of network equipment onto commodity hardware.

   This use case document provides resiliency requirements for
   virtualization of the LTE mobile core network, known as virtualized
   EPC (vEPC).

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

   Mobile operators have deploying Long Term Evolution (LTE) Evolved
   Packet Core (EPC) to meet the mobility, latency and bandwidth
   requirements for a variety of mobile data users. The EPC is the
   latest evolution of the [3GPP-R8] core network architecture, and is
   based on IP.

   The EPC architecture is said to have a "flat architecture" with

minimal components and functions. Principally the design is
intended to minimise the number of function nodes required
and protocol conversation of mobile data traffic. However, EPC
elements are bespoke stand-alone hardware (i.e., different boxes
for different functions). Network operators have identified that
this approach costly and inflexible.

The ETSI Network Functions Virtualization (NFV) Industry Steering
Group (ISG) published a set of use cases [NFV-ISG-UC]. One key
use case described the Virtualisation of Mobile Core Network and IP
Multimedia Subsystem (IMS), known as the vEPC.

The NFV approach takes the EPCs functional elements and runs them as
software instances (Virtual Appliances) on high-volume industry-
standard generic servers. This approach has number of advantages
including:

o Reducing: Cost, Power, Space and Complexity.

o Increasing: Flexibility, Scalability and Consolidation.

This use case document describes the vEPC architecture, functional
components and defines the resiliency requirements for the vEPC use
case.

1.1 Operator Benefits of Virtualization

There are a number of Operator Benefits which can be achieved
through virtualization of the EPC, these include:

o Economies of scale through common virtualized platform
o Enables a Multi-Service (MS) platform
o Reducing time to market to offer new services

o Uniformity of operations
o Simplified high availability
o Simplified disaster recovery
o Preferred test and diagnostic tools embedded

o Simplified in-service software upgrades
o Reduced training
o Simplified planning and provisioning
o Automation of installation
o Reduced site visits

2. Terminology

Evolved Packet Core (EPC): is an evolution of the 3GPP GPRS system

characterized by a higher-data-rate, lower-latency, packet-
optimized system.

Home Subscriber Server (HSS): a database that contains
user-related and subscriber-related information. It also provides
support functions in mobility management, call and session setup,
user authentication and access authorization.

Mobility Management Entity (MME) provides the signaling related
to mobility and security for Evolved UMTS Terrestrial Radio Access
Network (E-UTRAN) access.

Packet Data Network Gateway (PDN GW): is the point of interconnect
between the EPC and the external IP networks.

Policy and Charging Rules Function (PCRF): provides policy and
service control and the appropriate interfaces towards the mobile
charging and billing systems.

Serving GW (SGW): is the interconnect between the radio-side and the
EPC. The SGW serves the User Equipment (UE) by routing the incoming
and outgoing IP packets.

Virtualized Network Function (VNF): a VNF provides the same
functional behavior and interfaces as the equivalent network
function, but is deployed as software instances building on top of a
virtualization layer.

VNF Pool: a group of VNF instances providing the same network
function.

VNF Pool Element: a VNF instance inside a VNF pool.

VNF Pool Manager: an entity that manages a VNF pool, and interacts
with the service control entity to provide the network function.

VNF Set: a group of VNF instances that can be used to build network
services.

3. Virtual Evolved Packet Core (vEPC)

Deploying and operating mobile core network functions on
commodity hardware resources may provide significant network usage
efficiency and reductions in operational expenditure. Increased
automation would also accommodate scaling of voice and mobile data
demands.

The ETSI NFV use case [NFV-ISG-UC] describes requirements for

   server and packet gateways used for Packet Data Network
   (PDN) connections and IP Multimedia Subsystem (IMS) session (see
   Figure 1: Virtualized mobile core network and IMS).

   Typically mobile services are typically time dependent and may
   require a large number of computing resources in proportion to the
   number of users and/or service requests. Therefore it is
   desirable to scale them according to their specific computing
   requirements. The virtualization can be applied to the Evolved
   Packet Core (EPC) and the IMS to provide end to end
   service with service availability and resilience.

3.1 Mobile Core Network Components

   Within the mobile core network a number of nodes and
   specific functions are currently provided by dedicated hardware
   and software for mobile voice and data services, these are
   described in more detail in the following sub-sections.

3.1.1 Mobile Network Nodes

   The EPC is comprised of a variety of nodes, these include:

   o Mobility Management Entity (MME);

   o Serving Gateway (SGW);

   o Packet Data Network Gateway (PDN-GW);

   o Home Subscriber Server (HSS).

3.1.2 Mobile Network Functions

   The EPC provides a number of functions to manage mobile user traffic,
   these include:

   o Firewall (FW);

   o Policy Control (PC);

   o Network Address Translation (NAT);

   o Load Balancing (LB);

   o Deep Packet Inspection (DPI);

   o TCP Optimization of Traffic Flows;

   o HTTP Enrichment of Traffic Flows;

o Video Stream Optimization;

o Video Content Caching.

3.2 vEPC Resiliency Requirements

   When those virtualized service nodes(e.g., virtualized S/P-GW and
   IMS functions) are failed or overloaded, dynamic relocation of
   VNFs can be performed, the relocation of the managed
   sessions and/or connections must be accordingly managed. It also
   should be noted in [NFV-REL-REQ] that the traffic in the original
   VSN must be routed to the new location and it is desirable that
   the movement of the VSN is transparent to other VSN and or
   physical network entities such as client application on the UE.
   That is to say the other VSNs do not require to take any special
   action to this movement.

```
    +---------------+   +------------------------------+
    | vEPC          |   | vIMS                         |
    |               |   |                              |
    |   +--------+  |   |                  +---------+  |
    |   |        |  |   |                  |         |  |
    |   | vP/SGW +---+-+-|            +--+ vS-CSCF   |  |
    |   |        |  | | |            |   |         |  |
    |   +--------+  | | | +--------+ |   +---------+  |
    |Overload/Failure| |-+-|        +---| Overload/Failure |
    |               | | | | P-CSCF |                 |
    |               | ++++|        +++++             |
    |   +--------+  | + | +--------+  +  +---------+  |
    |   |        |  | + |           +  |         |  |
    |   | vP/SGW ++++++ |          +++| vS-CSCF  |  |
    |   |        |  | |  |           |         |  |
    |   +--------+  | |  |          +---------+  |
    |               | |  |                       |
    |  PDN Connection| |  | IMS Session          |
    +---------------+   +------------------------------+
```
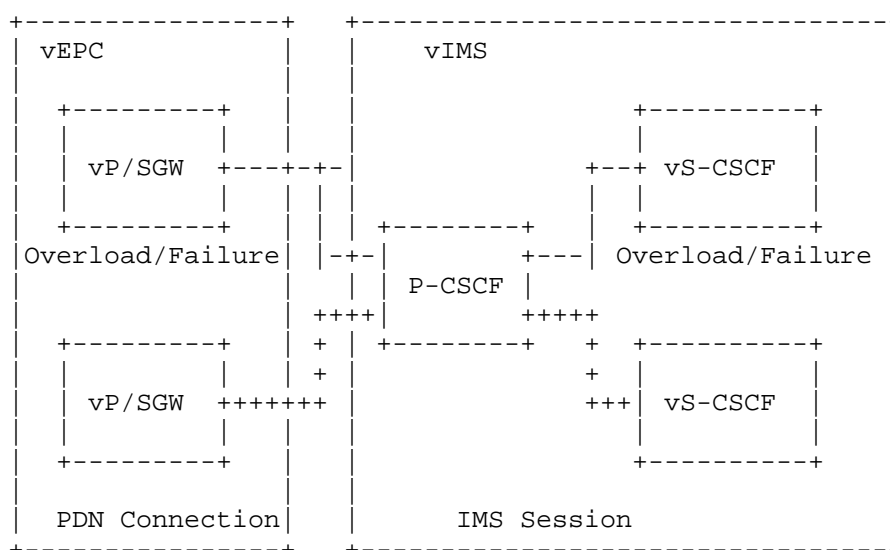
               Figure 1: Virtualized Mobile Core Network and IMS

   In this architecture, the following general resiliency requirements
   need to be satisfied:

   o Resource scaling - elastic service aware resource allocation to
      network functions;

   o State maintenance - network and network function state management
      during VSN relocation, replication, and resource scaling;

   o Monitoring/fault detection/diagnosis/recovery - appropriate

       mechanism for monitoring/fault detection/diagnosis/recovery of all
       components and their states after virtualization, e.g. VNF,
       hardware, hypervisor;

    o Service Availability - achieving the same level of service
       availability for the end-to-end virtualized mobile core network as
       in non-virtualized networks with reduced cost;

    o Minimum impact on other relevant functions.

3.2.1 Handling Unplanned Traffic Peaks

   Vendors are currently working with the Japanese Government to
   demonstrate the capabilities that a vEPC can have in handling
   unplanned traffic surges due to unforeseen circumstances:

   o A recent earthquake in Japan caused the demand for calls to
   increase to 150% capacity in the effected area. Calls were dropped
   due to the network capacity.

   o At the time the capacity in other areas was only 50%. In a vEPC
   environment the free resources from the other areas could have been
   used to manage this additional load.

3.2.2 Scaling of Resources and Functions

   The Evolved Packet System (EPS) is built from logical network
   functions, e.g. MME, PDN Gateway, Serving Gateway and Radio Base
   station (evolved NodeB) which are connected through the specified
   architectures references points. The 3GPP standard considers load
   balancing between different logical network functions of the same
   type. For example, Radio Base stations can choose one out of
   multiple available MMEs according to load-based weight factors to
   register an attaching mobile device. Mobile network operators can
   dimension their network in terms of numbers of required MMEs or
   data gateways according to statistical figures and thorough
   network planning, such as busy hour call attempts (BHCA).

   Virtualization technology enables adding additional resources as
   logical network functions by means of instantiation of the relevant
   functions in virtual machines. The instantiation of additional
   virtualized PDN Gateways or MMEs requires the announcement of their
   availability to other network components of the EPS. New attachments
   can then be balanced and distributed between an increased number of
   available network functions. Such procedure for scale-out suits the
   adaptation of the EPS resources to an increasing demand with low time
   constraints, e.g. due to an expected increase in subscribers or
   traffic volume.

    Unexpected increase in traffic or subscribers' attempt to request
    mobile service can result from scheduled events, e.g. festivals, or
    in particular after disaster events, such as an earthquake. The
    latter case in particular requires the mobile network to handle
    service requests and traffic from a huge amount of active mobile
    subscribers.

    Communication services during disaster events are essential, not only
    to provide a communication platform for rescue workers, but also to
    allow private subscribers to communicate with relatives.

    Such unexpected increase in active subscribers and traffic volume
    should not result in dropped connections, e.g. forced disconnects to
    offload existing subscriber states and traffic volume. It is
    preferable to scale-out resources internal to a single logical
    network function, e.g. an MME or a PDN Gateway. The advantage of
    such network function-internal resources scaling is the in-dependency
    of and transparency to external network functions and EPC protocols.

    Functionality and resources for a particular Virtualized Network
    Function (VNF) may be provisioned by the interplay of multiple
    virtualized Network Function Components (VNFC), whose instances
    map 1:1, or m:1, to virtual machines. Scaling up internally of a
    single instance of a VNF may be accomplished by the instantiation
    of additional VNFC instances. Load on the VNF must then be
    balanced between the multiple VNFC instances (LB). Such scaling
    must remain transparent to external network entities and to other
    VNFs.

                        Virtual Network Function

```
              +----------------------------+
              |    +--+          +----+     |
    <========>|LB|<--+--->|vNFC|     |
              |    +--+     |    +----+     |
              |  +----+     |    +----+     |
              |  |vNFC|<--+--->|vNFC|     |
              |  +----+          +----+     |
              |                            |
              +----------------------------+
```
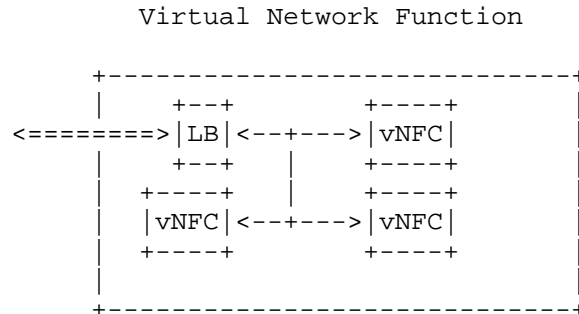
       Figure 2: Composition of multiple VNFC instances to build a single
              VNF.

    Technology for VNF scaling must also provide means to scale-in and
    reduce the number of resources in terms of required VNFCs, which
    provide the required network function.

Technology for VNF scaling must also provide means to scale-in and
reduce the number of resources in terms of required VNFs, which
provide the required network function.

Some general requirements for scaling in the view of virtualized EPC
network functions:

o Transparency and compatibility of network functions virtualization
to legacy EPS components;

o Support for scale-out of VNFs, representing additional logical
EPC network functions;

o Inter-working with configuration management (OSS) to configure and
announce new Network Functions to the EPS;

o Automation of scaling and simplified OAM;

o VNF-internal scale-out and resiliency management;

o Support of scale-in and associated shut down of VNFC instances;
handling of states associated with VNFCs, which are to be shut
down (state depletion vs. state transfer/offload);

o (non-critical: VM aggregation to fewer host servers, e.g. to enable
host server power saving).

Service requirements for the scaling of VNFs from VNFPool
perspective, based on the current working group scope of
work:

o Balancing load between VNFs within a VNFPool;

o Inter-working with system-wide (e.g. EPS) load balancing, e.g.
  cellular-specific selection of VNFs;

o Compatibility with system-wide addressing of selected VNFs.
  VNFPool solutions may consider different addressing schemes
  and associated address mapping within and outside a VNFPool;

o Coordination of scale-out and scale-in of VNFs within a VNFPool;

o Coordination of the use, visibility and addressability of
  additional VNF resources. New VNFs, which carry a new system-wide
  identifier, need to be announced to the system. New VNFs, which
  carry only a new VNFPool-internal identifier and provide additional
  VNF resources for an existing instance of a network function
  (system is aware of the network function instance's
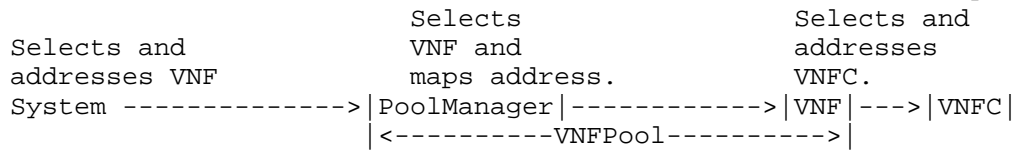 identifier) require only VNFPool-internal coordination.

```
                         Selects                    Selects and
      Selects and        VNF and                    addresses
      addresses VNF      maps address.              VNFC.
      System ------------->|PoolManager|------------->|VNF|--->|VNFC|
                           |<----------VNFPool---------->|
```

      Figure: Scope of VNFPool and coordination between VNFPool-internal
              and system-wide selection, balancing and addressing of
              network functions

3.2.3 Failure Handling

   During vEPC deployment, various failures can occur, for instance
   virtual machine failure, hypervisor failure, a broken host server,
   failure in a datacenter's transport network infrastructure, as well
   as failure of network links which connect a datacenter to the global
   network infrastructure.

   It is unlikely that a single solution suits the handling of all kind
   of failures. Typically for today's products, function redundancy and
   state synchronization as well as failure detection and failover are
   function and implementation specific.

   The detection of VM or hardware failures on a host server, as well as
   failure of networking equipment may introduce some delay before the
   system initiates failover to standby or backup resources. It may not
   be possible for an operator to meet agreed service levels in all
   cases.

   Due to the variety of different failure reasons, detection of the
   failure type may be required to initiate the appropriate procedure
   for failover handling. Mobile operators have strong requirements to
   minimize the time of system outage as experiences by subscribers,
   hence require minimal detection and failover handling latencies.

   Referring to the architecture of a virtualized Network Function as
   depicted in Figure 2, some VNFCs may require synchronization of
   states with a standby VNFC instance of the same kind to introduce
   redundancy on VNFC level. Others may not require state
   synchronization but rely simply a backup VNFC with the same
   functionality, as in case of failure, states can be recovered and
   retrieved from a different VNF, which holds the same or a sub-set of
   these states. Hence, redundancy management and failover mechanisms
   can be VNFC-specific.

   Disaster events, such as an earthquake, can have impact to the
   availability of a larger VNF Set (a group of VNFs providing different
   functions) or even to the access to a complete data center in case
   the data center's links to the global network infrastructure

breaks. In such case, even the availability of a backup system in
a globally and topologically distant data center can meet the
requirement of service continuation. Seamless continuation of
subscribers' services is unlikely, as it would require maintenance
of state synchronization between functions being instantiated in
different data centers. But solely the provisioning of backup vNFs
allows subscribers to re-attach to the mobile communication system
and place new calls. Handling such failover requires macroscopic
indirection of the EPC reference points to a set of backup VNFs in
a different data center.

Some general requirements for failure detection and failover handling
in the view of virtualized EPC network functions:

o Support function-specific redundancy and failover management;

o Support different kinds of redundancy for failover (state
synchronization between VNF instances, state recovery at backup VNF
instances, state re-establishment at a backup VNF instance);

o Selection of appropriate commodity hardware for backup and failover
(resources availability);

o Minimize state synchronization- and failover latency;

o Detection of failure;

o Detection of failure type and level (e.g. VNF, hypervisor,
hardware, network);

o Enforcement of failover strategy according to failure type;

o Automated detection and failure handling.

Service requirements for failure handling from VNFPool perspective,
based on the current working group scope of work:

o Selection of suitable resources (host server, rack, topological
  location) for redundant VNFs;

o Instantiation and installation of redundant resources on VNF-level;

o Policing and enforcement of different redundancy schemes (e.g.
  active/standby synchronization, backup VNF);

o Inter-working between VNF-internal (active/standby VNFC) and
  external (VNF redundancy) redundancy management;

o Failover between VNFs within a VNFPool;

   o Handling of VNFPool-internal addressing and identification in case
     of failover;

```
                Addresses
                VNF
         System ------>|Pool Manager|--+-->VNF-----+->VNFC
                                    |       failover|
                              failover,|            +->VNFC
                              address  |
                              handling +-->VNF
                       |<------VNFPool ----->|
```

     Figure: Scope of VNFPool and coordination between VNFPool-internal
             when handling failures.

3.2.4 State Synchronization

   vEPC components may be split into control (signaling) and forwarding
   (data) plane traffic. A failure of a control plane traffic may result
   in the loss of communication between EPC functions. This should not
   impact user forwarding traffic, and it may be necessary for control
   functions to have state maintained and synchronized with back-up
   VNF instances hosting control elements.

   Also it may be necessary for data plane state to also be
   synchronized so certain connections continue to be operational
   and capable of forwarding traffic during from one VNF to
   another.

3.3 What does that mean for Virtual Network Function Pool (VNF Pool)?

   For VNF Pool in the view of EPC, it is to be investigated where an
   IETF-based generalized functional architecture and common protocol
   can support vEPC scaling, failure detection and handling. Such
   common protocol components should allow inter-working with VNF-
   specific and possibly proprietary but highly efficient mechanisms
   for redundancy and fault management.

   The granularity of a VNF Pool Manager[zong-vnfpool-problem-
   statement] may be a VNF, VNF Pool or VNF Set. It is assumed that a
   Pool Manager handles VNFs with the granularity of EPC network
   functions (MME, PDN Gateway).

   A VNF Pool Manager's role for load balancing between PEs is to be
   investigated, taking additional and independent load balancing
   instances for macroscopic (system-wide) load balancing within the EPS
   and for microscopic load balancing (between multiple VNFs of a
   single logical VNF instance) into account.

3.3.1 VNF Pool Definitions

   There is a hierarchy of terms used to describe VNF Pool components
   and their relationship:

   o An instantiation of a VNF is known as a VNF instance;

   o A group of VNF instances is known as a VNF Set;

   o A managed VNF Set is known as a VNF Pool;

   o A VNF pool is managed using a VNF Pool Manager.

   These definitions will be moved into the terminology section if they
   are agreed by the working group.

4. IANA Considerations

   This document makes no IANA requests.

5. Security Considerations

   [To be discussed.]

6. References

6.1. Normative References

6.2. Informative References

   [3GPP-R8]

   [NFV-ISG-UC]
            "Network Function Virtualisation; Use Cases;", ISG NFV Use
            Case, June 2013.

   [NFV-REL-REQ]
            "Network Function Virtualisation Resiliency Requirements",
            ISG REL Requirements, June 2013.

   [zong-vnfpool-problem-statement]
            Zong, N., "Problem Statement for Reliable Virtualized
            Network Function (VNF) Pool", May 2014.

Authors' Addresses

   Peter Willis
   British Telecom
   UK

     Email: peter.j.willis@bt.com

     Daniel King
     Lancaster University
     UK

     Email: d.king@lancaster.ac.uk

     Jeong-dong Ryoo
     ETRI

     Email: ryoo@etri.re.kr

     Marco Liebsch
     NEC Laboratories Europe

     Email: liebsch@neclab.eu