

CUDB Virtual Infrastructure Requirements

USER GUIDE

Copyright

© Ericsson AB 2016, 2017. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



Contents

1	Introduction	1
1.1	Revision Information	1
2	Compute Requirements	3
3	Network Requirements	7
4	Storage Requirements	17
5	Security Requirements	19
6	Other Requirements	21
	Glossary	23
	Reference List	25





1 Introduction

This document describes the requirements for the infrastructure as requested by virtualized CUDB.

1.1 Revision Information

Rev. A

Initial release.

Rev. B

Other than editorial changes, this document has been revised as follows:

- Section 2 on page 3: Updated reference in Compute Host row of Table 1.

Rev. C

Other than editorial changes, this document has been revised as follows:

- Updated the Requirement Text throughout the document.
- Section 2 on page 3: Removed virtualized CPU row from Table 1. Added information about CPU Pinning, NUMA, and VM affinity or anti-affinity rules to Table 1.
- Section 3 on page 7: Removed IP Multicast from Table 2 and added information about Trunk virtualized NIC support.
- Section 5 on page 19: Removed Trunk virtualized NIC support row from Table 4.
- Section 6 on page 21: Removed Para-virtualized drivers row from Table 5 and added information about DPDK.

Rev. D

Other than editorial changes, this document has been revised as follows:

- Section 3 on page 7: Updated the category definition and requirement text of Routing protocol category in Table 2.





2 Compute Requirements

This section lists all compute requirements, see Table 1 for details.

Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Physical CPU architecture	<p>A physical CPU in its simplest terms refers to a physical CPU core, that is, a physical Hardware Execution Context (HEC), but can refer to a processor that is manufactured to contain multiple physical cores.</p> <p>If the physical CPU supports hyperthreading, then that enables a single processor core to act like two processors, that is, logical processors.</p> <p>[ETSI definition], Reference [4]: Device in the compute node, which provides the primary container interface. This is the generic processor that executes the code of the Virtualized Network Function Component (VNFC).]</p>	<p>Physical CPUs with x86_64 architecture in the host that also supports: VT-x/AMD-V hardware acceleration and hyper-threading technology.</p> <p>Hyper-threading is recommended to be enabled.</p> <p>Note: Validation and verification of virtualized CUDB was performed on single socket Generic Ericsson Processor version 5 (GEP5) boards that are equipped with Intel XEON E5-2658v2 (Ivy Bridge) processor.</p>
CPU pinning	<p>CPU pinning (or processing affinity) allows virtual CPUs (vCPUs) used by guests (Virtual Machine (VM) process or threads) to be bound to a physical CPU (pCPU). This is achieved by configuring the appropriate policies or parameters in the scheduler responsible for allocating computing resources.</p>	<p>CPU pinning is recommended so that it is controlled that each vCPU is allocated to a fixed pCPU (to prevent free float of vCPU across cores), and different vCPUs of the same VM are placed on the same pCPUs. The goal is that a VM does not share a pCPU with other VMs.</p> <p>Not using any CPU pinning or using a different CPU pinning policy than the one recommended above is also possible, though performance and predictability may be impacted by the final strategy.</p>



Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Number of virtualized CPUs	[ETSI definition] , Reference [4]: VM is a virtualized computation environment that behaves very much like a physical computer or server. A VM has all its ingredients (processor, memory/storage, interfaces/ports) of a physical computer or server and is generated by a hypervisor (see Section 6 on page 21), which partitions the underlying physical resources and allocates them to VMs. VMs are capable of hosting a VNFC.]	<p>The number of virtualized CPUs is determined at dimensioning activity.</p> <p>The minimum number of virtualized CPUs per VM is 2.</p>
NUMA	<p>NUMA (Non-Uniform Memory Access) is a hardware design that separates its cores into multiple clusters, where each cluster has its own local memory region and still allows cores from one cluster to access all memory in the system (access to own local memory is faster than access to non-local memory). Each cluster is called NUMA node, and it consists of processors and local memory.</p> <p>Any system that knows that it is running over a NUMA hardware so that it can make decisions (usually performance-related) based on the underlying NUMA topology is called NUMA aware.</p>	<p>Having a NUMA-aware scheduler for compute resources is recommended so that the vCPUs of a VM are always allocated inside a NUMA node. Even though spanning a VM over more than one NUMA node is possible, it is not recommended, as performance might be impacted due to increased latency.</p>



Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Memory	<p>Volatile RAM requires power to maintain the stored information. It retains its contents while powered on, but when the power is interrupted the stored data is lost very rapidly or immediately.</p> <p>[ETSI definition, Reference [4]: This represents the virtual memory needed for the Virtualization Deployment Unit (VDU) or the VM. The VDU is a construct used in an information model and the Virtualized Network Function (VNF) can be modeled using one or multiple such constructs, as applicable.]</p>	<p>The memory requirements are determined at the dimensioning activity.</p> <p>The minimum memory capacity required is 6 GB per VM.</p> <p>If available, it is advised to use Huge pages (typically, 1GB page size) for the memory allocation of virtualized CUDB VMs.</p>
Compute host	<p>A compute host (or simply host) is the whole server entity providing computing resources, composed of the underlying hardware platform: processor, memory, I/O devices, and disk. The hypervisor (see Section 6 on page 21) may or may not be seen as part of the host.</p> <p>[No ETSI definition]</p>	<p>The recommended minimum number of compute hosts with hardware and software redundancy is at least four.</p> <p>Providing the actual number of compute hosts requires a system dimensioning activity. The number of hosts and resources must fulfill the Virtual Deployment Considerations in the CUDB Deployment Guide.</p> <p>For further details, refer to <i>CUDB Deployment Guide</i>, Reference [1].</p>
VM affinity or anti-affinity rules	<p>The relationship among VMs and hosts to control the placement of VMs in the infrastructure for High Availability (HA) purposes.</p>	<p>It is required to follow affinity rules defined in <i>CUDB Deployment Guide</i>, Reference [1] and in <i>CUDB High Availability</i>, Reference [2].</p>

Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Overcommitting CPU	<p>CPU overcommitting is a hypervisor feature (see Section 6 on page 21) that allows a VM to allocate more virtualized CPUs than physical CPUs the host has available.</p> <p>The term overallocation is also used for this feature.</p> <p>[ETSI definition, Reference [5]: The VDU may coexist on a platform with multiple VDUs or VMs and is as such sharing CPU core resources available in the platform. It may be necessary to specify the CPU core oversubscription policy in terms of virtual cores to physical cores or threads on the platform. This policy can be based on required VDU deployment characteristics such as high performance, low latency, and deterministic behavior.]</p>	<p>Overcommitting CPU is not allowed.</p> <p>It compromises the predictability and dimensioning of capacity, latency, quality of service, and other characteristics of the VM.</p>
Overcommitting memory	<p>Memory overcommitting is a hypervisor feature (see Section 6 on page 21) that allows the sum of all VM memory allocations to be bigger than the total memory of the host.</p> <p>The term overallocation is also used for this feature.</p> <p>[No ETSI definition]</p>	<p>Overcommitting memory is not allowed.</p> <p>It compromises the predictability and dimensioning of capacity, latency, quality of service, and other characteristics of the VM.</p>



3 Network Requirements

This section lists all network requirements, see Table 2 for details.

Table 2 Network Requirements

Category	Category Definition	Requirement Text
Virtualized NICs per VM	[ETSI definition] , Reference [4]: <ul style="list-style-type: none"> • Network Interface Card (NIC) is a device in a compute node that provides a physical interface with the infrastructure network. • Virtualized NIC is created for a VM by a hypervisor.] 	Number of required virtualized NICs per VM depends on VM type: <ul style="list-style-type: none"> • System Controller (SC) VM: 6 • Payload VM: 7
Trunk virtualized NIC support	To support a high number of VLANs.	Trunk virtualized NIC support is not required.
Virtual networks or VLANs	A Virtual Local Area Network (VLAN) is the logical grouping of network nodes, which allows geographically dispersed network nodes to communicate as if they were physically on the same network. [ETSI definition] , Reference [6]: Virtual network is a topological component used to affect the forwarding of specific characteristic information. The virtual network is bounded by its set of permissible network interfaces. Virtual network forwards information among the network interfaces of VM instances and physical network interfaces, providing the necessary connectivity, and ensures the secure isolation of traffic from different virtual networks.]	Using VLANs is recommended for network isolation on the internal networks. The number of VLAN per VM depends on the VM type: <ul style="list-style-type: none"> • SC VM: 6 • Payload VM: 7

Table 2 Network Requirements

Category	Category Definition	Requirement Text
Bandwidth of internal network	<p>Internal network is a virtual network used for Transparent Inter-Process Communication (TIPC), Internal INET, and boot traffic.</p> <p>The bandwidth is measured on the virtualized NIC assigned to the internal network.</p>	<p>The requirement for bandwidth depends on the traffic and data models. Providing the actual bandwidth requirement involves a system dimensioning activity.</p> <p>The minimum values for standard VoLTE traffic following the Ericsson standard model are (Rx/Tx values) per VM:</p> <ul style="list-style-type: none"> • 105.50 MB/s / 112.09 MB/s
Bandwidth of external networks	<p>External networks are the virtual networks used for communication external to the VNF. For example, network function (other VNFs or PNFs), network management systems, and charging system.</p> <p>The bandwidth of external networks is the sum of the measured bandwidth of all virtualized NICs (except the virtualized NIC for VNF internal network) connected to the VMs in the VNF.</p>	<p>The requirement for bandwidth depends on the traffic and data models. Providing the actual bandwidth requirement involves a system dimensioning activity.</p> <p>The minimum values for a VNF servicing standard VoLTE traffic following the the Ericsson standard model are (Rx/Tx values)</p> <ul style="list-style-type: none"> • 35.35 MB/s / 27.49 MB/s
Pinning virtualized NICs	<p>Pinning virtualized NICs to physical ports enables to manage the distribution of traffic. When pinning is set, all traffic from the virtualized NIC travels through the I/O module to the specified Ethernet port.</p> <p>[No ETSI definition]</p>	<p>Pinning virtualized NICs to physical ports is not required.</p>
L2 redundancy	<p>To achieve telecom grade failure recovery, the virtualized NIC interface is protected in the L2 infrastructure, for example, by using two physical NICs to achieve resiliency in the external switches, in case one switch plane is broken (assuming duplicated L2 switch).</p> <p>[No ETSI definition]</p>	<p>Telecom grade availability of the virtual network is required, therefore L2 redundancy must be secured by the cloud infrastructure.</p>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
L2/L3 QoS	<p>Quality Of Service (QoS) settings at L2/L3 for the traffic are not changed within the virtual network boundaries.</p> <p>[ETSI definition, Reference [6]: Describes the QoS options to be supported on the Virtual Link (VL), for example, latency and jitter.]</p>	Differentiated Services Code Point (DSCP) passthrough is required.
L3 network separation	<p>Overlap between the IP addresses used for a given network, and the IP addresses used for part of another network, where these networks are adjacent in the communication path.</p> <p>[No ETSI definition]</p>	Virtual Routers (VRs) must be used per Traffic Type, typically to separate OAM from User signalling traffic. If physical separation is possible, those VRs must be assigned physically separated uplinks. Traffic separation serves also the purpose of a more secure network design.
Virtualized NIC type	<p>Virtualized NIC can be of access or trunk type. Each virtualized NIC can have multiple IP interfaces either of the same or of different type.</p> <p>IP aliasing is the concept of creating or configuring multiple IP addresses on a single network interface. In dual-stack configuration, the device is configured for both IPv4 and IPv6 network stacks. The dual-stack configuration can be implemented on a single interface or with multiple interfaces. In this configuration, the device decides how to send the traffic based on the destination address of the other device.</p> <p>[No ETSI definition]</p>	Access type virtualized NICs are required.

Table 2 *Network Requirements*

Category	Category Definition	Requirement Text
IP address allocation	<p>The process of assigning IP addresses to the virtualized NICs that are associated with the VNF, including the permission for the assigning.</p> <p>[No ETSI definition]</p>	It must be possible to self-assign addresses to the virtualized NIC instances.
Path supervision	<p>Any path supervision protocols can be used, such as Gratuitous Address Resolution Protocol (ARP), Internet Control Message Protocol (ICMP), or Bidirectional Forwarding Detection (BFD).</p> <p>[No ETSI definition]</p>	BFD support is required.
L3 redundancy	<p>L3 redundancy can be provided by the Virtual Router Redundancy Protocol (VRRP).</p> <p>[No ETSI definition]</p>	VRRP support is required.
Booting network	<p>The Preboot eXecution Environment (PXE) specification describes a standardized client-server environment that boots a software assembly, retrieved from a network, on PXE-enabled clients. On the client side, it requires only a PXE-capable NIC, and uses a small set of industry-standard network protocols, such as Dynamic Host Configuration Protocol (DHCP) and Trivial File Transfer Protocol (TFTP). The DHCP is a standardized network protocol used on IP networks for dynamically distributing network configuration parameters, such as IP addresses for interfaces and services.</p> <p>[No ETSI definition]</p>	The virtualization infrastructure must allow PXE booting.
IPv4 or IPv6	<p>Internet Protocol version 4 (IPv4) and 6 (IPv6).</p> <p>[No ETSI definition]</p>	IPv4 needs to be supported.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Routing protocol	<p>Open Shortest Path First (OSPF) is an Interior Gateway routing protocol for IP networks based on the shortest path first or link-state algorithm.</p> <p>BFD is a network protocol used to detect faults between two forwarding engines connected by a link, even on physical media that do not support failure detection of any kind. Static routing is a form of routing that occurs when a router uses a manually configured routing entry, rather than information from a dynamic routing traffic. Static routes are fixed and do not change if the network is changed or reconfigured.</p> <p>Equal-Cost Multipath (ECMP) is a routing strategy, where next-hop packet forwarding to a single destination can occur over multiple “best paths”, which tie for top place in the routing metric calculations.</p> <p>[No ETSI definition]</p>	For deployments using OSPF and route supervision between the VNFs and routers, OSPF and BFD capable routers are necessary. For deployments using static routing and no route supervision, OSPF and BFD capability is not necessary. ECMP capability is required in both options.
LBaaS	<p>LBaaS is a feature available through OpenStack Neutron. It allows for proprietary and open-source load balancing technologies to drive the actual load balancing of requests, allowing OpenStack operators to use a common interface and move seamlessly between different load balancing technologies.</p> <p>[No ETSI definition]</p>	No specific requirements apply.

*Table 2 Network Requirements*

Category	Category Definition	Requirement Text
NTP	<p>NTP is a networking protocol for clock synchronization between computer systems over packet-switched, variable-latency data networks.</p> <p>[No ETSI definition]</p>	All the VM instances must be able to access an appropriate NTP server.
DNS	<p>The Domain Name System (DNS) is a hierarchical distributed naming system for computers, services, or any resource connected to Internet or to a private network. It translates domain names, which can be easily memorized by humans, to the numerical IP addresses.</p> <p>[No ETSI definition]</p>	No specific requirements apply.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Latency	<p>Network latency in a packet switched network is measured either one way (the time from the source sending a packet to the destination receiving it), or round-trip delay time (the one-way latency from source to destination plus the one-way latency from the destination back to the source).</p> <p>For a definition, refer to ITU-T Y.1540, Reference [8] and ITU-T G.1020, Reference [9]. For the recommended values, refer to ITU-T Y.1541, Reference [10] and ITU-T G.114, Reference [11].</p> <p>[ETSI definition, Reference [7]: Packet delay is the elapsed time between a packet being presented to the Network Function Virtualization (NFV) virtual network from one VNFC guest OS instance to that same packet being presented to the destination VNFC guest OS instance. Packets that are delivered with more than the maximum acceptable packet delay for the VNF are counted as packet loss events and excluded from packet delay measurements.]</p>	NFV Infrastructure latency must be less than 500 microseconds.

Table 2 *Network Requirements*

Category	Category Definition	Requirement Text
Jitter	<p>In packet switched networks, jitter is the variation in latency as measured in the variability over time of the packet latency across a network. Packet jitter is expressed as an average of the deviation from the network mean latency.</p> <p>For a definition, refer to ITU-T Y.1540, Reference [8], ITU-T G.1020, Reference [9], and RFC 3393, Reference [12]. For the recommended values, refer to ITU-T Y.1541, Reference [10].</p> <p>[ETSI definition, Reference [7]: Packet delay variance (that is, jitter) is the variance in packet delay.]</p>	Jitter is required to meet general Telecom grade requirements.
Packet loss	<p>Packet loss occurs when one or more packets of data traveling across a computer network fail to reach their destination. Packet loss is measured as a percentage of packets lost divided by packets sent.</p> <p>For a definition, refer to ITU-T Y.1540, Reference [8] and to ITU-T G.1020, Reference [9]. For the recommended values, refer to ITU-T Y.1541, Reference [10].</p> <p>[ETSI definition, Reference [7]: Packet loss is the rate of packets that are either never delivered to the destination or delivered to the destination after the maximum acceptable packet delay of the VNF.]</p>	Packet loss for the NFV infrastructure must be less than 0.001%.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
VLAN tagging	<p>VLAN tagging is used to separate the traffic of different VLANs when the VLANs span multiple switches. VLAN tagging is done by inserting a VLAN ID into a packet header to identify to which VLAN the packet belongs.</p> <p>[No ETSI definition]</p>	The externally routed networks must use VLAN tagging.
MTU Size	<p>The Maximum Transmission Unit (MTU) is the largest packet size, measured in bytes that can be transmitted over a network. Any messages larger than the MTU are divided into smaller packets before being sent. Breaking them up slows down transmission speeds. Ideally, the MTU size must be the same as the smallest MTU size of all the networks between the local computer and a message's final destination.</p> <p>Fragmentation in IPv6 is performed only by source nodes, not by routers along a packet's delivery path.</p>	MTU must be set to 1500 bytes.
Virtualized NIC sub-interfacing		ARP unicast that contains MAC address, which is not assigned to the interfaces by the infrastructure should be allowed.





4 Storage Requirements

This section lists all storage requirements, see Table 3 for details.

Table 3 Storage Requirements

Category	Category Definition	Requirement Text
Storage	<p>Persistent storage space used for storing and retrieving digital information.</p> <p>[ETSI definition, Reference [5]: Required storage characteristics (for example, size), including Key Quality Indicators (KQIs) for performance and reliability/availability.]</p>	<p>The requirement for VM storage size depends on the hardware traffic and data models.</p> <p>Providing the actual storage capacity requirement involves a system dimensioning activity.</p> <p>The minimum storage needs are:</p> <ul style="list-style-type: none"> • System Controller VM: 80 GB • Payload VM: 20 GB
Storage performance	<p>Performance capability of a storage device is determined by the following three factors:</p> <ul style="list-style-type: none"> • Speed or throughput or bandwidth: the speed at which data is transferred out of or into the storage device (normally measured in megabytes per second). • Latency: how long it takes for a storage device to start an I/O task (measured in fractions of a second). • Speed values vary depending on the access operation (sequential or random). <p>[ETSI definition, Reference [6] for latency: The latency in accessing a specific state held in storage to execute an instruction cycle.]</p>	<p>The requirement for read and write speed depends on the traffic and data models.</p> <p>Providing the actual storage performance requirement involves a system dimensioning activity.</p> <p>The minimum values per VM for standard VoLTE traffic following the Ericsson standard model are (Read/write speed):</p> <ul style="list-style-type: none"> • 2.112 MB/s / 110.433 MB/s





5 Security Requirements

This section lists all security requirements, see Table 4 for details.

Table 4 Security Requirements

Category	Category Definition	Requirement Text
Virtualized NIC traffic separation	Different types of traffic are separated to provide security.	Network separation must be maintained
Virtual Switch traffic separation	Different types of traffic are separated to provide security.	See L3 network separation in Table 2.
Physical interfaces traffic separation	Different types of traffic are separated to provide security.	No specific requirements apply.
VNF isolation by the hypervisor	VNFs are to be protected and isolated from other VNFs in the environment.	The hypervisor must ensure the security of VNFs by preventing interferences from other VNFs in the deployment, that is, memory, storage, and other resources assigned to a VNF are not accessible from other VNFs.
Hypervisor security against VM escape attempts	VMs are protected and isolated from other VMs in the environment.	The hypervisor must prevent VNFs from escaping to the hypervisor. The hypervisor software is to be upgraded to remove security issues (several vulnerabilities on different hypervisors have been reported, which allows VNF to escape to the hypervisor).
OAM authentication and authorization	OAM protection of the hypervisor.	The hypervisor must implement proper authentication and authorization mechanisms to prevent unauthorized users from accessing the hypervisor and performing malicious activities. Different accounts with different roles must be implemented. Audit trails logs must be implemented.

Table 4 *Security Requirements*

Category	Category Definition	Requirement Text
OAM access control to VNFs	Restrict access to VNFs.	The hypervisor must implement control over which hypervisor accounts are capable of managing specific VNFs.
IP packet filtering	IP packet filtering functionality.	<p>Cloud Infrastructure routers and switches must provide tools to:</p> <ul style="list-style-type: none"> • Perform IP filtering and VLAN encapsulation. • Disable Processing of Packets Utilizing IP Options or alternative means. • Support filtering based on the value(s) of any portion of the protocol headers for IP, TCP, UDP, and ICMP. • Filter both incoming and outgoing traffic on any IP interface. The router to be used in cloud deployments must provide a similar feature to ACLs. • Provide a iptables-like feature for equivalent filtering of traffic from or to the VMs. • It must be possible to log all routing or traffic filter actions.



6 Other Requirements

This section lists all other requirements, see Table 5 for details.

Table 5 Other Requirements

Category	Category Definition	Requirement Text
Hypervisor	<p>A hypervisor, or Virtual Machine Monitor (VMM), is a piece of computer software, firmware, or hardware that creates and runs VMs. A computer, on which a hypervisor is running one or more VMs, is defined as a host machine. Each VM is called a guest machine. The hypervisor presents the guest operating systems with a virtual operating platform and manages the execution of the guest operating systems. Multiple instances of various operating systems can share the virtualized hardware resources.</p> <p>[ETSI definition, see Reference [6]: Hypervisor is a piece of software that partitions the underlying physical resources, and creates VMs and isolates the VMs from each other; running either directly on top of the hardware (bare metal hypervisor) or running on top of a hosting Operating System (OS) (hosted hypervisor). The abstraction of resources comprises all those entities inside a computer or server that are accessible, like processor, memory/storage, or NICs. The hypervisor enables the portability of VMs to different hardware.]</p>	<p>Virtualized CUDB is a software-only product verified with QEMU-KVM on x86 64-bit processors with VT-x extension. In theory, any kind of hypervisor can be suitable that meets the computing, virtual networking and storage-related infrastructure requirements. The hypervisor must support the SUSE Linux Enterprise Server 12 (SLESv12) guest OS.</p>

Table 5 Other Requirements

Category	Category Definition	Requirement Text
DPDK	DPDK (Data Plane Development Kit) is a set of libraries and drivers for fast packet processing. It is designed to run on any processor.	Using DPDK is recommended.
Installation	Any tools and environment-related software that is needed for installation.	Heat Orchestration Template (HOT) - version 2014-10-16 -based installation method must be supported. The HOT based installation for CEE/Openstack includes <code>qcow2</code> images.
Cloud administration related security		Cloud administrative operations must not simultaneously impact compute nodes hosting redundant virtualized CUDB components (refer to <i>CUDB High Availability</i> , Reference [2]) to not impact the VNF functionality.
Geographical distribution across datacenters		All elements of a CUDB VNF must be instantiated in the same datacenter.



Glossary

For the terms, definitions, acronyms and abbreviations used in this document, refer to *CUDB Glossary of Terms and Acronyms*, Reference [3].





Reference List

CUDB Documents

- [1] *CUDB Deployment Guide*
- [2] *CUDB High Availability*
- [3] *CUDB Glossary of Terms and Acronyms*

Other Documents and Online References

- [4] *Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV* http://www.etsi.org/deliver/etsi_gs/NFV/001_099/003/01.02.01_60/gs_NFV003v010201p.pdf
- [5] *Network Functions Virtualisation (NFV); Management and Orchestration* http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf
- [6] *Network Functions Virtualisation (NFV); Infrastructure Overview* http://www.etsi.org/deliver/etsi_gs/NFV-INF/001_099/001/01.01.01_60/gs_NFV-INF001v010101p.pdf
- [7] *Network Functions Virtualisation (NFV); Service Quality Metrics* http://www.etsi.org/deliver/etsi_gs/NFV-INF/001_099/010/01.01.01_60/gs_NFV-INF010v010101p.pdf
- [8] *Y.1540 : Internet protocol data communication service - IP packet transfer and availability performance parameters* <https://www.itu.int/rec/T-REC-Y.1540>
- [9] *G.1020 : Performance parameter definitions for quality of speech and other voiceband applications utilizing IP networks* <https://www.itu.int/rec/T-REC-G.1020/en>
- [10] *Y.1541 : Network performance objectives for IP-based services* <https://www.itu.int/rec/T-REC-Y.1541/en>
- [11] *G.114 : One-way transmission time* <https://www.itu.int/rec/T-REC-G.114/en>
- [12] *IP Packet Delay Variation Metric for IP Performance Metrics (IPPM). IETF RFC 3393* <https://www.ietf.org/rfc/rfc3393.txt>