

Virtual CSCF Infrastructure Requirements

Call Session Control Function

REQUIREMENTS SPECIFICATION

Copyright

© Ericsson AB 2016–2018. 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
2	Compute Requirements	3
3	Network Requirements	9
4	Storage Requirements	21
5	Security Requirements	23
6	Other Requirements	25





1 Introduction

This document describes the minimum infrastructure resource requirements to deploy the virtual Call Session Control Function (vCSCF) in a cloud deployment.

Note: The scope of this document is for a single Virtual Machine (VM) per host deployment. Multiple VMs running on each host is not covered in this document.





2 Compute Requirements

This section lists all compute requirements, see Table 1.

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. But it can also 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><i>[ETSI definition: Device in the compute node which provides the primary container interface. This is the generic processor, which executes the code of the VNF⁽¹⁾ Component.]</i></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>The identification of the virtual VNF infrastructure requirement was performed on GEP5⁽²⁾ boards that are equipped with Intel XEON E5-2658v2 (Ivy Bridge) processor.</p>
vCPU ⁽³⁾	<p><i>[ETSI definition: The vCPU created for a VM by a hypervisor (see Section 6 on page 25). In practice, a vCPU can be a time sharing of a real CPU or in the case of multicore CPUs, it can be an allocation of one or more cores to a VM.]</i></p> <p>vCPU-affinity can be used to isolate a physical CPU to a vCPU, by pinning the vCPU to a dedicated physical CPU.</p>	<p>vCPU affinity is recommended to be used when multiple VMs are present on the compute host.</p> <p>Use vCPU affinity to ensure that vCPUs of a VM never shares (threads within) a physical CPU core with vCPUs of other VMs.</p>



Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Number of vCPUs	<i>[ETSI definition: 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 25), which partitions the underlying physical resources and allocates them to VMs. VMs can host a VNFC]</i>	<p>The number of vCPUs depends on physical resources.</p> <p>The number of vCPUs for a PL⁽⁴⁾ VM is at least 2. The reference number is 8.</p> <p>The number of vCPUs for an SC⁽⁵⁾ VM is at least 2. The reference number is 8.</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 rapidly or immediately.</p> <p><i>[ETSI definition: This represents the virtual memory needed for the VDU⁽⁶⁾ or VM. VDU is a construct used in an information model and the VNF can be modeled using one or multiple such constructs, as applicable.]</i></p>	<p>The required minimum amount of memory of a PL⁽⁴⁾ VM with 2 vCPUs is 24 GB. The reference memory is 32 GB for 8 vCPUs and 48 GB for 16 vCPUs respectively.</p> <p>The required minimum amount of memory of a SC⁽⁵⁾ VM with 2 vCPUs is 8 GB. The reference memory is 16 GB for 8 vCPUs and 24 GB for 16vCPUs respectively.</p> <p>The memory use depends on the compute node hardware characteristics and traffic model of the network. The numbers are indicative of the memory use required for a VM running at an average of 50% of CPU on GEP5⁽⁷⁾ hardware, with a traffic mix consisting of:</p> <ul style="list-style-type: none"> • 12% - Originating B2B • 26% - Terminating B2B • 12% - PSTN • 50% - Third party registration <p>To determine the actual VM memory use, refer to <i>vCSCF Dimensioning and Characteristics Specification</i>.</p>

Table 1 Compute Requirements

Category	Category Definition	Requirement Text
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 25) can be seen as part of the host.</p> <p><i>[No ETSI definition]</i></p>	<p>The recommended minimum number of compute hosts with hardware and software redundancy is 4.</p> <p>Separate compute hosts are recommended for each of the 4 VMs (that is, 2 SC and 2 PL VMs).</p> <p>Hardware redundancy SC VMs belonging to the same VNF must not be collocated on the same compute host for redundancy reasons.</p> <p>PL VMs belonging to the same VNF are not recommended to be collocated on the same compute host for redundancy reasons.</p> <p>Support for VM anti-affinity policies or a solution to control the VM placement on compute hosts.</p>



Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Overcommitting CPU	<p>CPU overcommitting is a hypervisor feature (see Section 6 on page 25) 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><i>[ETSI definition: The VDU can coexist on a platform with multiple VDUs or VMs and as such is sharing CPU core resources available in the platform. It can be necessary to specify the CPU core oversubscription policy in terms of virtual cores to physical cores/threads on the platform. This policy can be based on required VDU deployment characteristics such as high performance, low latency, or deterministic behavior.]</i></p>	Overcommitting CPU is not allowed. It compromises the predictability and dimensioning of capacity, latency, quality of service and other characteristics of the VM.
Overcommitting memory	<p>Memory overcommitting is a hypervisor feature (see Section 6 on page 25) 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><i>[No ETSI definition]</i></p>	Overcommitting memory is not allowed. It compromises the predictability and dimensioning of capacity, latency, quality of service and other characteristics of the VM.

(1) Virtualized Network Function (VNF)

(2) Generic Ericsson Processor 5 (GEP5)

(3) Virtual CPU (vCPU)

(4) Payload (PL)

(5) System Controller (SC)

(6) Virtualization Deployment Unit (VDU)

(7) GEP5-64 boards with XEON E5-2658v2 (Ivy Bridge) processor and 64 GB memory





3 Network Requirements

This section lists the recommended network requirements, see Table 2.

Table 2 Network Requirements

Category	Category Definition	Requirement Text
vNICs ⁽¹⁾ per VM	<p><i>[ETSI definition: NIC is a device in a compute node that provides a physical interface with the infrastructure network]</i></p> <p><i>[ETSI definition: vNIC is a virtualized NIC created for a VM by a hypervisor.]</i></p>	<p>An SC VM requires 2 vNICs.</p> <p>A PL VM requires 4 vNICs.</p> <p>vNICs must be presented to the VM in a deterministic order during boot, since different networks and functions are statically assigned to the vNIC based on the order as they appear in the VM.</p>
Virtual networks or VLANs ⁽²⁾ per vNIC	<p><i>[ETSI definition: Virtual network is a topological component used to affect forwarding of specific characteristic information.</i></p> <p><i>The virtual network is bounded by its set of permissible network interfaces.</i></p> <p><i>Virtual network forwards information among the network interfaces of VM instances and physical network interfaces, providing the necessary connectivity and ensures secure isolation of traffic from different virtual networks.]</i></p> <p>A 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.</p>	<p>The following 4 VLAN separated virtual networks are required each by the CSCF having its own access port vNIC:</p> <ul style="list-style-type: none"> • VNF Internal VLAN • O&M VLAN • Signalling VLAN • Charging VLAN • Confidential VLAN <p>The virtual networks VLANs are not visible to the VNF, that is, traffic is not tagged.</p>

Table 2 Network Requirements

Category	Category Definition	Requirement Text
Bandwidth of the internal network	<p>Internal network is a virtual network used for TIPC, Internal INET, and boot traffic.</p> <p>The bandwidth is measured on the vNIC assigned to the internal network.</p>	<p>The virtualization infrastructure must provide at least 20 Mbps for a 2+2 VNF with 8 vCPU VM.</p> <p>The virtualization infrastructure must provide at least 120 Mbps for a 2+6 VNF with 12 vCPU VM.</p> <p>The bandwidth of the VNF internal network has impact on the duration of, for example, the synchronization between the block storage on SCs (during installation), scaling of the VNF, and VM/VNF reboot.</p> <p>The bandwidth use depends on the compute node hardware characteristics and traffic model of the network. The numbers are indicative of the bandwidth on the internal network required for a VM running at an average of 50% CPU on GEP5⁽³⁾ hardware, with a traffic mix consisting of:</p> <ul style="list-style-type: none"> • 12% - Originating B2B • 26% - Terminating B2B • 12% - PSTN • 50%- Third party registration <p>To determine the actual internal network bandwidth, refer to <i>vCSCF Dimensioning and Characteristics Specification</i>.</p>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Bandwidth of the external virtual networks	<p>External networks are the virtual networks used for communication external to the VNF. For example, network functions (other VNFs or PNFs⁽⁴⁾), Network Management Systems, charging system.</p> <p>The sum of the measured bandwidth of all vNICs connected to the VM.</p>	<p>The virtualization infrastructure must provide at least 20 Mbps per PL and 5 Mbps per SC.</p> <p>For an SC, the O&M virtual network normally uses < 5 Mbps of the bandwidth. The O&M virtual network includes different types of communication such as SNMP⁽⁵⁾, NETCONF, file transfer (PM⁽⁶⁾, backups, software downloads).</p> <p>The detailed bandwidth use per PL is as follows:</p> <ul style="list-style-type: none"> • Charging virtual network ~ 2 Mbps • Signalling virtual network ~ 15 Mbps <p>The bandwidth use depends on the compute node hardware characteristics and traffic model of the network. The numbers are indicative of the bandwidth of the external network required for a VM running at an average of 50% of CPU on GEP5 (16) hardware, with a traffic mix consisting of:</p> <ul style="list-style-type: none"> • 12% - Originating B2B • 26% - Terminating B2B • 12% - PSTN • 50%- Third party registration <p>To determine the actual bandwidth for external networks, refer to <i>vCSCF Dimensioning and Characteristics Specification</i>.</p>

Table 2 Network Requirements

Category	Category Definition	Requirement Text
Pinning vNICs	<p>Pinning vNICs to physical ports enables managing the distribution of traffic. When pinning is set, all traffic from the vNIC travels through the I/O module to the specified Ethernet port.</p> <p><i>[No ETSI definition]</i></p>	Pinning vNICs to physical ports is not required.
L2 redundancy	<p>To achieve telecom grade failure recovery, the vNIC interface is protected in the L2 infrastructure, for example, by using two physical NICs to achieve resiliency in the external switches, if one switch plane is broken (assuming duplicated L2 switch).</p> <p><i>[No ETSI definition]</i></p>	Telecom grade availability of the virtual network is required for the CSCF, therefore L2 redundancy must be secured by the virtualization infrastructure.
L2/L3 QoS ⁽⁷⁾	<p>QoS settings at L2/L3 for the traffic are not changed within the virtual network boundaries.</p> <p>[ETSI definition: Describes the QoS options to be supported on the VL, for example, latency and jitter.]</p>	This Quality of Service (QoS) setting must be preserved end-to-end between the VNF and the next node or boundaries of the VLAN.
L3 network separation	<p>Overlap between the IP addresses used for a given network, and the IP addresses used for part of the other network, where these networks are adjacent in the communication path.</p> <p><i>[No ETSI definition]</i></p>	No specific requirements apply.
L2 path diversity	<p>Having multiple routes at L2 to reach a destination.</p> <p><i>[No ETSI definition]</i></p>	No specific requirements apply.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
vNIC type	<p>vNIC can be of access or trunk type. Each vNIC can have multiple IP interfaces either of the same or different type.</p> <p>IP aliasing is the concept of creating or configuring multiple IP addresses on a single network interface.</p> <p>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><i>[No ETSI definition]</i></p>	<p>Access port vNICs are required.</p> <p>Support of multiple IP interfaces on a vNIC is provided but not required.</p>
IP address allocation	<p>The process of assigning IP addresses to the vNICs that are associated to the VNF, including the permission to be able to make that assignment.</p> <p><i>[No ETSI definition]</i></p>	<p>The CSCF must be able to create its own IP interfaces. The virtualization infrastructure can assign subnets to the CSCF as long as the IP addresses in these subnets can be used freely by the CSCF.</p> <p>NAT or NAPT cannot be used in the virtualization infrastructure.</p>

Table 2 Network Requirements

Category	Category Definition	Requirement Text
Path supervision	Any path supervision protocol can be used, like Gratuitous ARP ⁽⁸⁾ , ICMP ⁽⁹⁾ , or BFD ⁽¹⁰⁾ . <i>[No ETSI definition]</i>	The virtualization infrastructure must support one of the following combinations: <ul style="list-style-type: none"> • Static Routing with BFD (recommended) • Static routing without BFD • OSPFv2/OSPFv3 with BFD • OSPFv2/OSPFv3 with short OSPF timer Gratuitous ARP and ICMP are required.
L3 redundancy	L3 redundancy can be provided by the VRRP ⁽¹¹⁾ . <i>[No ETSI definition]</i>	VRRP support is required when static routing without BFD is used for traffic networks.
Booting network	The PXE ⁽¹²⁾ 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 DHCP ⁽¹³⁾ and TFTP ⁽¹⁴⁾ . 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. <i>[No ETSI definition]</i>	The virtualization infrastructure must allow PXE booting and DHCP traffic generated on the CSCF internal virtual network. The CSCF provides an internal DHCP service.
IPv4 or IPv6	Internet Protocol version 4 (IPv4) and 6 (IPv6).	Virtualization infrastructure must support IPv4 or IPv6 at the transport layer.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Routing protocol	<p>OSPF⁽¹⁵⁾ 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.</p> <p>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><i>[No ETSI definition]</i></p>	<p>The virtualization infrastructure must support one of the following combinations:</p> <ul style="list-style-type: none"> • Static Routing with BFD (recommended) • Static routing without BFD • OSPFv2/OSPFv3 with BFD • OSPFv2/OSPFv3 with short OSPF timer <p>The virtualization infrastructure must support flow-based ECMP⁽¹⁶⁾ routing.</p>
LBaaS ⁽¹⁷⁾	<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><i>[No ETSI definition]</i></p>	No specific requirements apply.
NTP ⁽¹⁸⁾	<p>NTP is a networking protocol for clock synchronization between computer systems over packet-switched, variable-latency data networks.</p> <p><i>[No ETSI definition]</i></p>	All VM instances must be able to access an appropriate NTP server. Clock synchronization from the host (or hypervisor) to guest VM must not be used.

Table 2 Network Requirements

Category	Category Definition	Requirement Text
DNS ⁽¹⁹⁾	<p>The DNS is a hierarchical distributed naming system for computers, services, or any resource connected to the Internet or a private network. It translates domain names, which can be easily memorized by humans, to the numerical IP addresses.</p> <p><i>[No ETSI definition]</i></p>	All VM instances must be able to access an appropriate DNS server.
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 and ITU-T G.1020.</p> <p>For recommended values, refer to ITU-T Y.1541 and ITU-T G.114.</p> <p><i>[ETSI definition: Packet delay is the elapsed time between a packet being presented to the NFV⁽²⁰⁾ 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.]⁽²¹⁾</i></p>	Latency is required to meet general Telecom grade requirements.



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, ITU-T G.1020, and RFC 3393.</p> <p>For recommended values, refer to ITU-T Y.1541.</p> <p><i>[ETSI definition: Packet delay variance (or jitter) is the variance in packet delay.]</i></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 and ITU-T G.1020.</p> <p>For recommended values, refer to ITU-T Y.1541.</p> <p><i>[ETSI definition: 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.]</i></p>	Packet loss is required to meet general Telecom grade requirements.

Table 2 Network Requirements

Category	Category Definition	Requirement Text
VLAN Tagging	<p>VLAN Tagging is used to separate the traffic of different VLANs when 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><i>[No ETSI definition]</i></p>	<p>No specific requirements apply.</p> <p>When using VLAN network separation, VLAN tagging is performed by the vSwitch and other physical network infrastructure, but is not visible to the VNF. For example, the CSCF sees all packets as not tagged by the different access vNICs.</p>
MTU Size	<p>MTU⁽²²⁾ 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 is the same as the smallest MTU size of all the networks between the local computer and the final destination of a message.</p> <p>Fragmentation in IPv6 is performed only by source nodes, not by routers along the delivery path of a packet.</p>	<p>The recommended default IP MTU setting for the CSCF is 1452 on the core network interfaces. This leaves 48 bytes for eVIP IPv6 internal tunneling.</p> <p>All VNFs support path MTU discovery for IPv4 and IPv6.</p>



Table 2 Network Requirements

- (1) Virtualized Network Interface Controller (vNIC)
- (2) Virtual Local Area Network (VLAN)
- (3) GEP5-64 boards with XEON E5-2658v2 (Ivy Bridge) processor and 64 GB memory
- (4) Physical Network Function (PNF)
- (5) Simple Network Management Protocol (SNMP)
- (6) Performance Management (PM)
- (7) Quality of Service (QoS)
- (8) Address Resolution Protocol (ARP)
- (9) Internet Control Message Protocol (ICMP)
- (10) Bidirectional Forwarding Detection (BFD)
- (11) Virtual Router Redundancy Protocol (VRRP)
- (12) Preboot eXecution Environment (PXE)
- (13) Dynamic Host Configuration Protocol (DHCP)
- (14) Trivial File Transfer Protocol (TFTP)
- (15) Open Shortest Path First (OSPF)
- (16) Equal-Cost Multipath (ECMP)
- (17) Load-Balancing-as-a-Service (LBaaS)
- (18) Network Time Protocol (NTP)
- (19) Domain Name System (DNS)
- (20) Network Function Virtualization (NFV)
- (21) There are other types of latencies defined in the [ETSI specification](#).
- (22) Maximum Transmission Unit (MTU)





4 Storage Requirements

This section lists all storage requirements, see Table 3.

Table 3 Storage Requirements

Category	Category Definition	Requirement Text
Storage	<p>Persistent storage space used for storing and retrieving digital information.</p> <p><i>[ETSI definition: Required storage characteristics (for example, size), including KQIs⁽¹⁾ for performance and reliability/availability.]</i></p>	<p>Each SC VM is recommended to be configured with a disk of 250 GB⁽²⁾.</p> <p>Additional storage space can be required depending on the amount of charging data that needs to be stored for charging backup handler or ACR storage, that is, Rf communications failure.</p>
Storage performance	<p>Performance capability of a storage device is determined by the following 3 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 (measured normally in megabytes per second) • IOPS: Input/Output Operations per Second (read and write) • Latency: how long it takes for a storage device to start an I/O task (measured in fractions of a second) <p>Speed and IOPS values vary depending on the access operation (sequential or random).</p> <p><i>[ETSI definition for latency: The latency in accessing a specific state held in storage to execute an instruction cycle.]</i></p>	<p>The required storage performance depends on the compute node hardware characteristics and traffic model of the network.</p> <p>Typical read and write speed of block storage is 40 Mbps that means the storage performance must provide at least 500 IOPS per VM (SC VM storage). The numbers are indicative of storage performance for a VM running at an average of 50% of CPU on GEP5⁽³⁾ hardware, with a traffic mix consisting of the following:</p> <ul style="list-style-type: none"> • 12% - Originating B2B • 26% - Terminating B2B • 12% - PSTN • 50%- Third-party registration⁽⁴⁾ <p>In order for read and write speed of block storage to increase to 120 Mbps, such as during upgrade. The storage performance must provide at least 1950 IOPS per VM (SC VM storage).</p> <p>If 16 vCore is used instead of 8 vCore, the storage performance must at least provide 900 IOPS for 40 Mbps and 3250 IOPS for 120 Mbps with the previously mentioned traffic mix.</p> <p>During initial installation, the read and write speed depends on the virtualization infrastructure. (SC VM block storage synchronization).</p> <p>To determine the actual required storage performance, refer to <i>vCSCF Dimensioning and Characteristics Specification</i>.</p>

(1) Key Quality Indicator (KQI)

(2) Example disk storage distribution: Disk size 250 GB = Swap (4 GB) + Boot (4 GB) + Log (20 GB) + other (222 GB). Log storage is 20 GB regardless of the disk size.

(3) GEP5-64 boards with XEON E5-2658v2 (Ivy Bridge) processor and 64-GB memory

(4) The block size is around 10 KB in GEP5. For information on CPU load dimensioning, refer to *vCSCF Dimensioning and Characteristics Specification*.





5 Security Requirements

This section lists all security requirements, see Table 4.

Table 4 Security Requirements

Category	Category Definition	Requirement Text
vNIC traffic separation	Different types of traffic are separated to provide security.	Traffic separation is secured in-line with Ericsson IMS security principles. For further details, refer to <i>CSCF VNF Network Connectivity Overview</i> .
Trunk vNIC support	To support a high number of VLANs.	Trunk vNICs are not used by the virtual CSCF. For further details, refer to <i>CSCF VNF Network Connectivity Overview</i> .
Virtual Switch traffic separation	Different types of traffic are separated to provide security.	Virtual Switches in the hypervisor must be capable of switching packets based on the VLAN tags and provide separation for traffic with different VLAN tags.
Physical interfaces traffic separation	Different types of traffic are separated to provide security.	No hard requirement on physical Separation. Traffic separation to be sorted out with VLAN segmentation on L2 level.
VNF isolation by the hypervisor	VNFs are protected/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 to be accessible to other VNFs.

*Table 4 Security Requirements*

Category	Category Definition	Requirement Text
Hypervisor security against VNF escape attempts	VNFs are protected and isolated from other VNFs 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 users from accessing the hypervisor and perform malicious activities. Different accounts with different roles must be implemented. Audit trails logs must be implemented.



6 Other Requirements

This section lists all other requirements, see Table 5.

Table 5 Other Requirements

Category	Category Definition	Requirement Text
Hypervisor	<p>A hypervisor, or VMM⁽¹⁾ is a piece of computer software, firmware, or hardware that creates and runs VMs.</p> <p>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><i>[ETSI: Hypervisor is a piece of software that partitions the underlying physical resources and creates VMs, and isolates the VMs from each other.</i></p> <p><i>The hypervisor is a piece of software running either directly on top of the hardware (bare metal hypervisor) or running on top of a hosting Operating System (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.]</i></p>	<p>The CSCF is a software-only product verified with qemu-KVM on X86_64 processors with VT-x extension.</p> <p>In theory, any kind of hypervisor can be suitable that meets the computing, virtual networking, and storage-related cloud requirements.</p> <p>The hypervisor must support Linux[®] distribution guest Operating System.</p>

Table 5 Other Requirements

Category	Category Definition	Requirement Text
Para-virtualized drivers	<p>Para-virtualization is a virtualization technique that presents a software interface to VMs that is similar, but not identical to, that of the underlying hardware. The intent of the modified interface is to reduce the portion of the execution time spent by the guest performing operations that are substantially more difficult to run in a virtual environment compared to a non-virtualized environment.</p> <p>Para-virtualized drivers are I/O device drivers that interact directly with the virtualization platform (with no emulation) to deliver disk and network access efficiently. This allows the disk and network subsystems to operate at near native speeds even in a virtualized environment, without requiring changes to existing guest Operating Systems.</p> <p><i>[No ETSI definition]</i></p>	<p>The CSCF requires support for para-virtualized drivers, for example, virtio for KVM or vmxnet3 for VMware.</p> <p>Para-virtualized drivers (significantly) improve the performance and capacity of the VMs.</p> <p>or</p> <p>Para-virtualized drivers are required to achieve high performance and capacity characteristics in the VMs.</p>
Installation	Any tools and environment-related software that is needed for installation.	The CSCF provides HOT based installation for CEE (OpenStack Kilo ⁽²⁾) and OVF 1.1 based installation for VMware.

(1) Virtual Machine Monitor (VMM)

(2) The HOT based installation for a vanilla OpenStack Kilo deployment can require changes to the supplied HOT templates