

Virtual MTAS Infrastructure Requirements

MTAS

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	19
5	Security Requirements	21
6	Other Requirements	23





1 Introduction

This document describes the minimum infrastructure resource requirements to deploy virtualized MTAS application in a cloud deployment.





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 (HEC), but can refer to a processor that 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: Device in the compute node, which provides the primary container interface. This is the generic processor, which executes the code of the VNFC⁽¹⁾.]</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>Pinned vCPUs are to belong to one and the same NUMA node, to avoid performance degradations. CPU governors are configured in performance mode with fixed CPU frequency and maximum CPU C-State set to 0, to ensure that processors perform predictably, in accordance with dimensioning guidelines.</p> <p>The identification of the virtual MTAS infrastructure requirement was performed on Generic Ericsson Processor 5 (GEP5) boards, which are equipped with Intel XEON E5-2658v2 (Ivy Bridge) processor and on GEP7 boards, which are equipped with Intel XEON E5-2658v4.</p>
vCPU ⁽²⁾	<p>[ETSI definition: The vCPU created for a VM⁽³⁾ by a hypervisor (see Section 6 on page 23). In practice, a vCPU can be a time sharing of a real CPU and/or in the case of multicore CPUs, it may be an allocation of one or more cores to a VM.]</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 should be used to ensure that vCPUs of a VM never share (threads within) a physical CPU core with vCPUs of other VMs.</p>



Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Number of vCPUs	[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 23), which partitions the underlying physical resources and allocates them to VMs. VMs are capable of hosting a VNFC.]	The required minimum number of vCPUs per VM is 3. The recommended minimum number of vCPUs per VM is 4.



Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Memory	<p>Volatile RAM⁽⁴⁾ 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>[ETSI definition: This represents the virtual memory needed for the VDU⁽⁵⁾ 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.]</p>	<p>Minimum memory budget required for Guest VM is to be considered divided into the following budgets:</p> <ul style="list-style-type: none"> • Budget for Guest OS: 13 GB at least, considering 2 GB reserved for MMAS and 11 GB for OS and CBA components. • Budget for for vDicos memory (“MultiMMap”): approximately 0.8 GB/vCPU, used for Application processes. This figure fulfills application memory requirements for all supported HW Concepts up to GEP7. • Budget for VM shared memory (used for the in-memory distributed and replicated database used by the application): roughly equal to Budget for vDicos memory. <p>These budgets are supposed to be assigned on the lowest higher GB - boundary, as huge pages VM memory management is strongly recommended, for performance reasons.</p> <p>For a better VM memory dimensioning, that takes into account HW concept-specific Capacity factors per vCPU, traffic profile, services provisioned, number of users, and other dimensioning parameters, it is recommended for the Ericsson staff to consult the internal product dimensioning guidelines and tools.</p> <p>As an example, a 4 vCPU's vMTAS Payload VM is at least 13 GB + 2 * (0.8 * 4) ~ 20 GB.</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 23) 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>Separate compute hosts are recommended for each of the four VMs (that is, 2 SC and 2 PL VMs).</p> <p>If persistent block storage service (such as Cinder) is not provided by the cloud environment, ephemeral storage should be used for the disk resources associated with the SC VMs. Thus, the SC VMs should run on server blades with sufficient disk capacity.</p>



Table 1 Compute Requirements

Category	Category Definition	Requirement Text
Overcommitting CPU	<p>CPU overcommitting is a hypervisor feature (see Section 6 on page 23) 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: The VDU can coexist on a platform with multiple VDUs or VMs and is as such 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, and/or 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 23) 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>

(1) Virtualized Network Function Component (VNFC)

(2) Virtual CPU (vCPU)

(3) Virtual Machine (VM)

(4) Random-Access Memory (RAM)

(5) Virtualization Deployment Unit (VDU)





3 Network Requirements

This section lists all network requirements, see Table 2.

Table 2 Network Requirements

Category	Category Definition	Requirement Text
vNICs ⁽¹⁾ per VM	<p>[ETSI definition: NIC is a device in a compute node that provides a physical interface with the infrastructure network.]</p> <p>[ETSI definition: vNIC is a virtualized NIC created for a VM by a hypervisor.]</p>	<p>An MTAS VM requires 4 vNICs with Virtio/vmxnet3 drivers:</p> <ul style="list-style-type: none">• eth0 - <MTAS VNF>_internal network• eth1 - <MTAS VNF>_OM_int network• eth2 - <MTAS VNF>_om_sc_sp network• eth3 - <MTAS VNF>_sig_sc_sp network <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>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Virtual networks or VLANs ⁽²⁾ per vNIC	<p>[ETSI definition: Virtual network is a topological component used to affect forwarding of specific characteristic information.</p> <p>The virtual network is bounded by its set of permissible network interfaces.</p> <p>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.]</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 four VLAN separated virtual networks are required by MTAS each having its own vNIC:</p> <ul style="list-style-type: none">• <MTAS VNF>_internal• <MTAS VNF>_OM_int• <MTAS VNF>_om_sc_sp• <MTAS VNF>_sig_sc_sp <p>It is recommended to use VLAN segmentation to separate virtual network.</p>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Bandwidth of 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 cloud infrastructure must provide at least around 40 Mb/s per VM of bandwidth for a 2+2 MTAS VNF with 4 vCPU VMs.</p> <ul style="list-style-type: none"> • During PL/cluster reboot, the throughput can increase to 200 Mb/s <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 usage is dependent on the compute node hardware characteristics and traffic model of the network. The numbers mentioned above are indicative of the bandwidth on the internal network required for a VM with VoLTE traffic model on GEP5 hardware. Use the characteristics and dimensioning documentation to determine the actual internal network bandwidth.</p>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Bandwidth of the total virtual networks	The sum of the measured bandwidth of all vNICs connected to the VM.	<p>The cloud infrastructure must provide at least the following bandwidth per VM for:</p> <ul style="list-style-type: none">• vNIC eth1 <<1 Mb/s• vNIC eth2 typically < 8 Mb/s (but it depends on: 1. How fast the charging data buffer is supposed to be emptied in the case of Rf link reestablishment, 2. How fast the charging data is supposed to be fetched from the ACR storage).• vNIC eth3 8 Mb/s for a 2+2 cluster with 4 vCPUs (around 64 Mb/s for a 2+10 cluster with 9 vCPUs) <p>Total throughput of 57 Mb/s non-blocking switching is required from the underlying virtual switching infrastructure per VM for a 2+2 cluster with 4vCPU.</p> <p>The bandwidth usage is dependent on the compute node hardware characteristics and traffic model of the network. The numbers mentioned above are indicative of the bandwidth of the external network required for a VM with VoLTE traffic on GEP5 hardware. Use the characteristics and dimensioning documentation to determine the actual bandwidth for external networks.</p>
Pinning vNICs	<p>Pinning vNICs to physical ports enables to manage 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>[No ETSI definition]</p>	Pinning vNICs to physical ports is not required.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
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, in case one switch plane is broken (assuming duplicated L2 switch).</p> <p>[No ETSI definition]</p>	Telecom grade availability (4×9's) of the virtual network is required for MTAS. Therefore, L2 redundancy must be secured by the virtualization infrastructure.
L2/L3 QoS ⁽³⁾	<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>	MTAS supports DSCP marking of outgoing packets therefore it is recommended that the Differentiated Services Code Points are considered at the L3 boundaries of the virtual networks.
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>	L3 network separation is not required for MTAS on the <MTAS VNF>_internal and <MTAS VNF>_OM_int networks. On the other hand, globally routable subnets have to be reserved for the OM (sysmgmt) and VIP FEE networks.
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>[No ETSI definition]</p>	<p>Support of multiple IP interfaces on access vNICs is provided. This includes the following:</p> <ul style="list-style-type: none"> Multiple IP interfaces of the same IP version (alias interfaces) are required on the <MTAS VNF>_internal network (NFS, CP-M MIP addresses). Dual stack configuration, that is, both IPv4 and IPv6 interfaces.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
IP address allocation	<p>The process of assigning IP addresses to the vNICs that are associated to the VNF, including the permission for the assigning.</p> <p>[No ETSI definition]</p>	<p>The MTAS application must be able to create its own IP interfaces. The cloud manager can assign subnets to the MTAS as long as the IP addresses in these subnets can be used freely by the MTAS application.</p> <p>Moreover the virtualization infrastructure must allow packets to pass through virtual ports regardless of the subnet associated with the network.</p>
Path supervision	<p>Any path supervision protocols can be used, such as Gratuitous ARP⁽⁴⁾, ICMP⁽⁵⁾, or BFD⁽⁶⁾.</p> <p>[No ETSI definition]</p>	<p>The virtualization infrastructure should support one of the following combinations:</p> <ol style="list-style-type: none">1. Static Routing with BFD (preferred)2. OSPFv2/OSPFv3 with BFD3. OSPFv2/OSPFv3 with short OSPF timer. Gratuitous ARP and ICMP is required. <p>For further details, refer to MTAS Internal and External Connectivity.</p>
L3 redundancy	<p>L3 redundancy can be provided by the VRRP⁽⁷⁾.</p> <p>[No ETSI definition]</p>	<p>Redundant L3 routers (gateway routers) are required to provide the redundant static default routes on the VIP FEE networks.</p> <p>Redundant gateways must provide a VRRP protected default route on the OM_ext infrastructure.</p>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Booting network	<p>The PXE⁽⁸⁾ 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⁽⁹⁾ and TFTP⁽¹⁰⁾.</p> <p>The Dynamic Host Configuration Protocol (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 not block PXE booting traffic generated by the MTAS. The cloud infrastructure provided DHCP service must be disabled.
IPv4 or IPv6	<p>Internet Protocol version 4 (IPv4) and 6 (IPv6).</p> <p>[No ETSI definition]</p>	Virtualization infrastructure must support IPv4/IPv6 at the transport layer.
Routing protocol	<p>OSPF⁽¹¹⁾ 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>[No ETSI definition]</p>	<p>The virtualization infrastructure should support one of the following combinations:</p> <ol style="list-style-type: none"> 1. Static Routing with BFD (recommended) 2. OSPFv2/OSPFv3 with BFD 3. OSPFv2/OSPFv3 with short OSPF timer. <p>The virtualization infrastructure shall support flow-based Equal-Cost Multipath routing (ECMP).</p> <p>For further details, refer to MTAS Internal and External Connectivity.</p>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
NTP ⁽¹²⁾	<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 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>	All the 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 the recommended values, refer to ITU-T Y.1541 and ITU-T G.114.</p> <p>[ETSI definition: Packet delay is the elapsed time between a packet being presented to the NFV⁽¹³⁾ 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.]⁽¹⁴⁾</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 the recommended values, refer to ITU-T Y.1541.</p> <p>[ETSI definition: 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 and ITU-T G.1020.</p> <p>For the recommended values, refer to ITU-T Y.1541.</p> <p>[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.]</p>	Packet loss is required to meet general Telecom grade requirements.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
VLAN tagging	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. [No ETSI definition]	No specific requirements apply. 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, MTAS sees all packets as 'untagged' by the different access vNICs).
MTU Size	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 should be the same as the smallest MTU size of all the networks between the local computer and a message's final destination. Note: Fragmentation in IPv6 is performed only by source nodes, not by routers along a packet's delivery path.	The recommended default IP MTU setting for IMSv VNF applications is 1452 on the core network interfaces.

- (1) Virtualized Network Interface Controller (vNIC)
- (2) Virtual Local Area Network (VLAN)
- (3) Quality of Service (QoS)
- (4) Address Resolution Protocol (ARP)
- (5) Internet Control Message Protocol (ICMP)
- (6) Bidirectional Forwarding Detection (BFD)
- (7) Virtual Router Redundancy Protocol (VRRP)
- (8) Preboot eXecution Environment (PXE)
- (9) Dynamic Host Configuration Protocol (DHCP)
- (10) Trivial File Transfer Protocol (TFTP)
- (11) Open Shortest Path First (OSPF)
- (12) Network Time Protocol (NTP)
- (13) Network Function Virtualization
- (14) There are other types of latencies defined in the ETSI specification.



4 Storage Requirements

This section lists all storage requirements, see Table 3.

Table 3 Storage Requirements



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: Required storage characteristics (for example, size), including KQIs⁽¹⁾ for performance and reliability/availability.]</p>	<p>Each System Controller (SC) VM is recommended to be configured with a disk of at least 200 GB. (The minimum is 50 GB but it can only be used for lab trials.)</p> <p>The chosen storage size of an SC VM must be able to hold all charging data (from all PL VMs) in the charging backup handler (ACR storage) in case of Rf communication failures.</p> <p>The use of persistent block storage (Cinder volume) as SC disk resources is recommended however if no persistent block storage service is provided by the cloud environment ephemeral storage can be used for the disk resources associated with the SCs.</p>
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)• 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>[ETSI definition for latency: The latency in accessing a specific state held in storage to execute an instruction cycle.]</p>	<p>Typical read and write speed of block storage is 40 Mb/s which can increase to 120 Mb/s during upgrade.</p> <p>The required storage performance is dependent on the compute node hardware characteristics and traffic model of the network. The numbers mentioned above are indicative of storage performance for a VM with VoLTE traffic model on GEP5 hardware. Use the characteristics and dimensioning documentation to determine the actual required storage performance.</p>

(1) Key Quality Indicator (KQI)



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 MTAS Internal and External Connectivity.
Trunk vNIC support	To support a high number of VLANs.	Trunk vNICs are not used by the virtualized MTAS for further details refer to MTAS Internal and External Connectivity.
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	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 to be accessible to other VNFs.
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).



Table 4 Security Requirements

Category	Category Definition	Requirement Text
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.
OAM access control to VNFs	Restrict access to VNFs.	OAM access control is required.



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. 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: Hypervisor is a piece of software that partitions the underlying physical resources and creates VMs, and isolates the VMs from each other.</p> <p>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.]</p>	<p>MTAS is 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 minimum cloud requirements as detailed above.</p> <p>The hypervisor shall support SLES 12 (SuSE Linux Enterprise Server) 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, the underlying hardware. The intent of the modified interface is to reduce the portion of the execution time spent for 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. 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>[No ETSI definition]</p>	MTAS requires support for para-virtualized drivers. This is required to achieve high performance and capacity characteristics in the system.
Installation	Any tools and environment-related software that is needed for installation.	In an Openstack environment, Heat Orchestration Template (HOT) based deployment method shall be used.

(1) Virtual Machine Monitor (VMM)