

# Virtualized IPWorks Infrastructure Requirements

---

## REQUIREMENTS SPECIFICATION

## **Copyright**

© Ericsson AB 2017, 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

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Compute Requirements</b>	<b>3</b>
<b>3</b>	<b>Network Requirements</b>	<b>7</b>
<b>4</b>	<b>Storage Requirements</b>	<b>17</b>
<b>5</b>	<b>Security Requirements</b>	<b>19</b>
<b>6</b>	<b>Other Requirements</b>	<b>21</b>





# 1 Introduction

This document describes the requirements for the infrastructure resource as requested by IPWorks virtualized applications.





## 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><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 VNFC<sup>(1)</sup>.]</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 if hyper-threading is supported.</p> <p>The identification of the virtualized VNF infrastructure requirement was performed on the different boards corresponding to IPWorks deployment configuration:</p> <p>For IPWorks 2+2 deployment:</p> <ul style="list-style-type: none"> <li>• GEP5 boards equipped with Intel XEON E5-2658v2 (Ivy Bridge) processor</li> <li>• GEP7L boards that are equipped with Intel XEON E5-2618Lv4 (Broadwell-EP) processor</li> </ul>
vCPU <sup>(2)</sup>	<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><i>[ETSI definition: The vCPU created for a VM<sup>(3)</sup> by a hypervisor (see Section 6 on page 21). In practice, a vCPU may be a time sharing of a real CPU and/or in the case of multi-core CPUs, it may be an allocation of one or more cores to a VM.]</i></p>	<p>vCPU affinity is recommended to be used when multiple VMs are present on the compute host.</p> <p>vCPU affinity should be used 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 21), which partitions the underlying physical resources and allocates them to VMs. VMs are capable of hosting a VNFC.]</i>	<p>For PL in 2+2 Deployment,</p> <ul style="list-style-type: none"> <li>• For DNS, ENUM, AAA, the requirement amount of vCPU per VM is 2~14 (for example, 2+2 Standard, 14).</li> <li>• For DHCP, the requirement amount of vCPU per VM is 2~8.</li> </ul> <p>For SC in 2+2 Deployment,</p> <ul style="list-style-type: none"> <li>• For DNS, ENUM, AAA, the requirement amount of vCPU per VM is 2~14 (for example, 2+2 Standard, 14).</li> <li>• For DHCP, the requirement amount of vCPU per VM is 2~8 (for example, 2+2 Standard, 8).</li> </ul>
Memory	<p>Volatile RAM<sup>(4)</sup> 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><i>[ETSI definition: This represents the virtual memory needed for the VDU<sup>(5)</sup> or VM. VDU is a construct used in an information model and the VNF can be modelled using one or multiple such constructs, as applicable.]</i></p>	<p>For PL in 2+2 Deployment,</p> <ul style="list-style-type: none"> <li>• For DNS, ENUM, AAA, the requirement amount of memory per VM is 8 GB.</li> <li>• For DHCP, the requirement amount of memory per VM is 16 GB.</li> </ul> <p>For SC in 2+2 Deployment,</p> <ul style="list-style-type: none"> <li>• For DNS, ENUM, AAA, the requirement amount of memory per VM is 40 GB.</li> <li>• For DHCP, the requirement amount of memory per VM is 8 GB.</li> </ul>
Compute host	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.	Recommend that VMs are distributed on different compute hosts. (especially for the same VM type).



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><i>[ETSI definition: 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/threads on the platform. This policy can be based on required VDU deployment characteristics such as high performance, low latency, and/or deterministic behavior.]</i></p>	<p>CPU Overcommitting 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>Memory overcommitting 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 <sup>(1)</sup> 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>SC-VM and PL-VM require the same number of vNICs: 3 (three) each SC-VM:</p> <ul style="list-style-type: none"> <li>• Internal (TIPC, NFS, TFTP, boot)</li> <li>• O&amp;M</li> <li>• Provisioning</li> </ul> <p>PL-VM:</p> <ul style="list-style-type: none"> <li>• Internal (TIPC, NFS, TFTP, boot)</li> <li>• Data</li> <li>• Traffic</li> </ul> <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 <sup>(2)</sup> per vNIC	<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><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>The following 3 (three) VLAN separated virtual networks are required each by IPWorks having its own vNIC:</p> <ul style="list-style-type: none"><li>• O&amp;M VLAN</li><li>• Provisioning VLAN</li><li>• Traffic VLAN</li></ul>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Bandwidth of internal network	<p>Internal network is a virtual network used for TIPC, NFS, TFTP, and boot traffic.</p> <p>The bandwidth is measured on the vNIC assigned to the internal network.</p>	<p>For a 2+2 Standard cluster, the bandwidth is about ~60 Mbps.</p> <p>Note: 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 provided here are indicative of the bandwidth on the internal network required for a VM running at an average of 50% CPU on 2+2 Standard Deployment: GEP5 (14)/GEP7L (14) hardware; on 2+2 Flexible Deployment: GEP5 (2~14)/GEP7L (2~14) hardware, with a traffic mix used normally in UDM and IPWorks. The characteristics and dimensioning documentation shall be used to determine the actual internal network bandwidth.</p>
Bandwidth of the total virtual networks	The sum of the measured bandwidth of all vNICs connected to the VM.	<p>The virtualization infrastructure must provide at least ~30 Mbps per VM of external (total) signaling bandwidth for IPWorks.</p> <p>The virtualization infrastructure must provide at least ~300 Mbps of external OAM bandwidth for IPWorks.</p>
Pinning vNICs	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.	Pinning vNICs to physical ports is not required for IPWorks.

Table 2 Network Requirements

Category	Category Definition	Requirement Text
L2 redundancy	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).	Telecom grade availability of the virtual network is required for IPWorks therefore L2 redundancy must be secured by the cloud infrastructure.
L2/L3 QoS <sup>(3)</sup>	QoS settings at L2/L3 for the traffic are not changed within the virtual network boundaries.  <i>[ETSI definition: Describes the QoS options to be supported on the VL, for example, latency and jitter.]</i>	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	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.	L3 network separation is not required for IPWorks on the internal networks. L3 network separation is required for external network.  Moreover, unique IP ranges must be secured for the VIP FEE networks.
L2 path diversity	Having multiple routes at L2 to reach a destination.	L2 path diversity is not required for IPWorks but L2 path diversity for SCTP multi-homing networks (VLANs) could be used in case it is supported by the cloud infrastructure and required by the customer.



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>Multiple IP interfaces on access vNICs is supported. This includes the following:</p> <ul style="list-style-type: none"> <li>Multiple IP interfaces of the same IP version (alias interfaces) are required on the <code>&lt;IPWVNF&gt;_internal_network</code> (NFS, TFTP, MIP addresses) and <code>&lt;IPWVNF&gt;_oam_network</code>.</li> <li>Dual stack configuration, that is, both IPv4 and IPv6 interfaces are supported.</li> </ul>
IP address allocation	The process of assigning IP addresses to the vNICs that are associated to the VNF, including the permission for the assigning.	<p>IPWorks must be able to create its own IP interfaces. The virtualization infrastructure can assign subnets to the IPWorks as long as the IP addresses in these subnets can be used freely by the IPWorks application.</p> <p>Moreover, the cloud infrastructure must allow packets to pass through virtual ports regardless of the subnet associated with the network.</p> <p>NAT nor NAPT cannot be used in the virtualization infrastructure.</p>
L3 redundancy	L3 redundancy can be provided by the VRRP <sup>(4)</sup> .	eVIP functionality in IPWorks VNF needs L3 redundancy.



Table 2 Network Requirements

Category	Category Definition	Requirement Text
Booting network	<p>The PXE<sup>(5)</sup> 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<sup>(6)</sup> and TFTP<sup>(7)</sup>.</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>The virtualization infrastructure must allow PXE booting and DHCP traffic generated on the IPWorks INTERNAL VLAN mentioned before.</p> <p>The IPWorks VNF provides an internal DHCP service.</p>
IPv4 or IPv6	Internet Protocol version 4 (IPv4) and 6 (IPv6).	Virtualization infrastructure must support both IPv4 and IPv6 at the transport layer.
Routing protocol	<p>OSPF<sup>(8)</sup> 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>A recommendation is that the virtualization infrastructure supports static routing.</p> <p>The virtualization infrastructure shall support flow based Equal-Cost Multi-Path routing (ECMP).</p>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
LBaaS <sup>(9)</sup>	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.	No specific requirements apply. LBaaS is not required.
NTP <sup>(10)</sup>	NTP is a networking protocol for clock synchronization between computer systems over packet-switched, variable-latency data networks.	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.
DNS	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.	All VM instances must be able to access an appropriate DNS server.

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 <a href="#">ITU-T Y.1540</a> and <a href="#">ITU-T G.1020</a>.</p> <p>For the recommended values, refer to <a href="#">ITU-T Y.1541</a> and <a href="#">ITU-T G.114</a>.</p> <p><i>[ETSI definition: Packet delay is the elapsed time between a packet being presented to the NFV<sup>(11)</sup> 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.]<sup>(12)</sup></i></p>	<p>Latency is required to meet general Telco-grade requirement, refer to <a href="#">ITU-T Y.1541</a> and <a href="#">ITU-T G.114</a>.</p>
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 <a href="#">ITU-T Y.1540</a>, <a href="#">ITU-T G.1020</a>, and <a href="#">RFC 3393</a>.</p> <p>For the recommended values, refer to <a href="#">ITU-T Y.1541</a>.</p> <p><i>[ETSI definition: Packet delay variance (that is, jitter) is the variance in packet delay.]</i></p>	<p>Jitter is required to meet general Telecom grade requirements.</p>



Table 2 Network Requirements

Category	Category Definition	Requirement Text
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 <a href="#">ITU-T Y.1540</a> and <a href="#">ITU-T G.1020</a>.</p> <p>For the recommended values, refer to <a href="#">ITU-T Y.1541</a>.</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	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 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. (e.g. the IPWorks sees all packets as 'untagged' by the different access vNICs).  VLANs must be tagged by the vSwitch. The vNICs must be configured to untag packets.
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 IPWorks is 1452 on the core network interfaces. This leaves 48 bytes for eVIP IPv6 internal tunneling.  All VNFs support path MTU discovery for IPv4 and IPv6.  All VNFs support path MTU discovery for IPv4 and IPv6. However, ICMP packets may not be allowed in access networks due to security concerns.

- (1) Virtualized Network Interface Controller (vNIC)
- (2) Virtual Local Area Network (VLAN)
- (3) Quality of Service (QoS)
- (4) Virtual Router Redundancy Protocol (VRRP)
- (5) Preboot eXecution Environment (PXE)
- (6) Dynamic Host Configuration Protocol (DHCP)
- (7) Trivial File Transfer Protocol (TFTP)
- (8) Open Shortest Path First (OSPF)
- (9) Load-Balancing-as-a-Service (LBaaS)
- (10) Network Time Protocol (NTP)
- (11) Network Function Virtualization
- (12) 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

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<sup>(1)</sup> for performance and reliability/availability.]</i></p>	<p>For 2+2 Deployment, each SC VM is recommended to be configured with a disk of 75~280 GB.</p> <p>PL VM has no requirement for disk.</p> <p>Additional storage space may be required depending on the amount of logs, backup handler and other OAM data.</p>
Storage performance	<p>Performance capability of a storage device is determined by the following three factors:</p> <ul style="list-style-type: none"> <li>• 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)</li> <li>• IOPS: Input/Output Operations per Second (read and write)</li> <li>• Latency: how long it takes for a storage device to start an I/O task (measured in fractions of a second).</li> </ul> <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 storage performance must provide at least 500 IOPS per VM for IPWorks (SC VM storage).</p> <p>Typical read and write speed of block storage is 40 Mb/s, which can increase to 130 Mb/s during upgrade.</p> <p>During initial installation the read and write speed depends on the virtualization infrastructure. (SC VM block storage synchronization).</p> <p>The required storage performance is dependent on the compute node hardware characteristics and traffic model of the network.</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.
Trunk vNIC support	To support a high number of VLANs.	Trunk port vNICs are not used.
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. Support for virtual switch traffic separation is required.
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	<p>VNFs are to be protected and isolated from other VNFs in the environment.</p> <p>VNFs are to be protected and isolated from other VNFs in the environment. VNFs shall not be allowed to access network, virtual memory, disk, CPU or other resources of other VNFs.</p>	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 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 unauthorized 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 (that is, system management) access control is required.
Deployment-related security	Applications can have more deployment requirements in the security area.	No specific requirements apply.



## 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<sup>(1)</sup>, 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><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>IPWorks is a software-only product verified with CEE on BSP (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 shall support the LDE which is based on SUSE Linux Enterprise Server 12 (SLESv12) 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 allow 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>IPWorks requires support for one of the following para-virtualized drivers:</p> <ul style="list-style-type: none"><li>• Virtio for KVM</li></ul> <p>This support is required to achieve high performance and capacity characteristics in the system.</p>
Installation	<p>Any tools and environment-related software that is needed for installation.</p>	<p>Heat Orchestration Template (HOT) based installation method is supported.</p> <p>The HOT based installation for CEE/Openstack includes qcow2 images.</p>

(1) Virtual Machine Monitor (VMM)