

# Single Server System Dimensioning Guide, CEE R6

## Cloud Execution Environment

---

### CONFIGURATION MODEL

**Copyright**

© Ericsson AB 2016. 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>  |
| 1.1      | Target Group                              | 1         |
| 1.2      | System Characteristics                    | 1         |
| <b>2</b> | <b>CEE System</b>                         | <b>2</b>  |
| <b>3</b> | <b>HW Requirements</b>                    | <b>2</b>  |
| 3.1      | Server Configuration                      | 4         |
| 3.2      | Network Configuration                     | 4         |
| 3.3      | CPU Configuration                         | 4         |
| 3.3.1    | Allocating Single CPUs for OVS            | 6         |
| 3.4      | RAM Configuration                         | 6         |
| 3.4.1    | Introduction                              | 6         |
| 3.4.2    | Configuration                             | 7         |
| 3.5      | Storage Configuration                     | 8         |
| 3.5.1    | Local Storage Disk Space                  | 8         |
| 3.5.2    | Disk Requirements for Atlas               | 10        |
| 3.5.3    | Disk Requirements for Nova Snapshots      | 10        |
| <b>4</b> | <b>Characteristics</b>                    | <b>11</b> |
| 4.1      | General System Limits                     | 11        |
| 4.2      | Orchestration Interface                   | 11        |
| 4.3      | Tenant Execution Environment              | 12        |
| 4.3.1    | Performance                               | 12        |
| 4.3.2    | Resiliency                                | 12        |
| 4.4      | Network                                   | 12        |
| 4.4.1    | Performance                               | 12        |
| 4.4.2    | Resiliency                                | 15        |
| 4.4.3    | Tenant Network Limitations                | 15        |
| 4.5      | Storage Limitations                       | 15        |
| 4.6      | In-Service Performance                    | 16        |
| <b>5</b> | <b>System Limitations</b>                 | <b>16</b> |
| 5.1      | OpenStack Deviations                      | 16        |
| 5.2      | SW Configurations and Options             | 17        |
| 5.2.1    | Number of Parallel Root Volume Operations | 17        |
| 5.3      | Not Supported                             | 17        |
| 5.4      | Limitations and Workarounds               | 17        |



|     |                       |           |
|-----|-----------------------|-----------|
| 5.5 | Update Limitations    | 17        |
|     | <b>Reference List</b> | <b>18</b> |



# 1 Introduction

This document describes the characteristics of Cloud Execution Environment (CEE) to enable dimensioning and understanding the limitations of CEE. It also describes requirements on HW combinations required for running CEE. The application can have additional requirements.

Storage is measured in gibibyte (GiB), tebibyte (TiB), and mebibyte (MiB) in this document.

1 GiB is equivalent to 1.074 GB.

The following words are used in this document with the meaning specified below:

|                   |  |
|-------------------|--|
| <b>vNIC</b>       | A virtual network interface card (vNIC) provides connectivity between the Cloud SDN Switch (CSS) and a VM. A configuration can provide several vNICs to a VM.  |
| <b>Interface</b>  | A network interface. Can be either a physical NIC (PHY) providing CSS with board external connectivity, or a virtual NIC connecting CSS to a VM.   |
| <b>PMD thread</b> | CSS uses a Poll Mode Driver (PMD) technique that continuously polls incoming packets from the NICs, that is, interrupts are not used. To be able to reliably handle all incoming packets, a software continuously polls the NIC queues for packets to be handled. This software is executing in one or more threads that are called PMD threads. The execution environment for the PMD threads is isolated from the Linux scheduler to be able to reliably handle a high sustained packet flow without interrupts or delays caused by being scheduled out. |
| <b>NUMA</b>       | Non-Uniform Memory Architecture  |

## 1.1 Target Group

Cloud Infrastructure providers and application designers.

## 1.2 System Characteristics

The characteristic features of CEE are the following:

- Support for one physical server
- OpenStack® SW deployment (for deviations, see Section 5.1 on page 16)



For information on system limitations, see Section 5 on page 16.

## 2 CEE System

This section summarizes the Single Server CEE system.

- The Single Server CEE operates on one Dell R630 server.
- The used switches, border gateways (BGWs), and firewalls are not part of the CEE product, these are the user's responsibility.
- External storage is not supported.
- vFuel is used from an installation laptop that is connected to the system for the time of CEE installation, update, and other maintenance activities. When the activity has been completed, vFuel is switched off by using *vFuel On Demand Use*.

## 3 HW Requirements

This section describes generic hardware and firmware requirements for CEE, based on the certified Ericsson CEE R6 HW. The actual configuration used by a specific installation needs to be adjusted to the needs of the application based on the recommendation below.

HW deployment outside the already certified ones requires system integration work.

The CEE Compute server connected to the environment provided by the customer is shown in Figure 1:




Figure 1 CEE HW Environment

The supported HW is listed below:

Table 1 Supported HW

| Item                       | Status    |
|----------------------------|-----------|
| <b>Server</b>              |           |
| Dell R630                  | Certified |
| <b>Networking</b>          |           |
| Customer-configured switch | -         |
| <b>Storage</b>             |           |
| Raid 0 for local disk      | Supported |
| Raid 1 for local disk      | Certified |



## 3.1 Server Configuration

This section describes the HW requirements for the server.

*Table 2 HW Configuration Example, Server*

| Aspect               | Requirement  |
|----------------------|--|
| CPU                  | 2x Intel® Xeon® Processor E5-2680 v3 (12 cores per processor, 2 HyperThreads (HT) per core), 48 HTs are available                    |
| RAM                  | 128 GiB or more, up to maximum capacity  |
| NIC                  | 1x10 GE Intel Niantic and 1x1 GE unspecified (Control)   |
| Onboard Disk         | At least 1000 GiB SSD or HDD: <ul style="list-style-type: none"><li>• 2x500 GiB for RAID 0</li><li>• 2x1000 GiB for RAID 1</li></ul> |
| Management Interface | Dell iDRAC   |

## 3.2 Network Configuration

This section describes HW configuration for networking.

The physical host network contains two interfaces:

- 1x1 GE NIC for the control network
- 1x10 GE NIC for the tenant traffic network

## 3.3 CPU Configuration

Refer to the *Configuration File Guide* for more information about the configuration procedure.

Since the system has limited CPU resources, it is necessary to assign the CPUs manually to the different resource owners in order to achieve optimal performance. The cores available for VMs are divided in two pools:

- One pool where VMs are floating and CPUs are over-committed
- And one pool with dedicated CPUs

These pools must be dimensioned according to the application needs.





The number of available CPU IDs depends on the CPU model. For example, the CPU allocation recommended for EPC is shown in Table 3 and Figure 2. This allocation is valid for Dell R630 with Intel Xeon E5-2680 v3 processor.

**Table 3 CPU Allocation for Dell R630 with Intel Xeon E5-2680 v3 Processor**

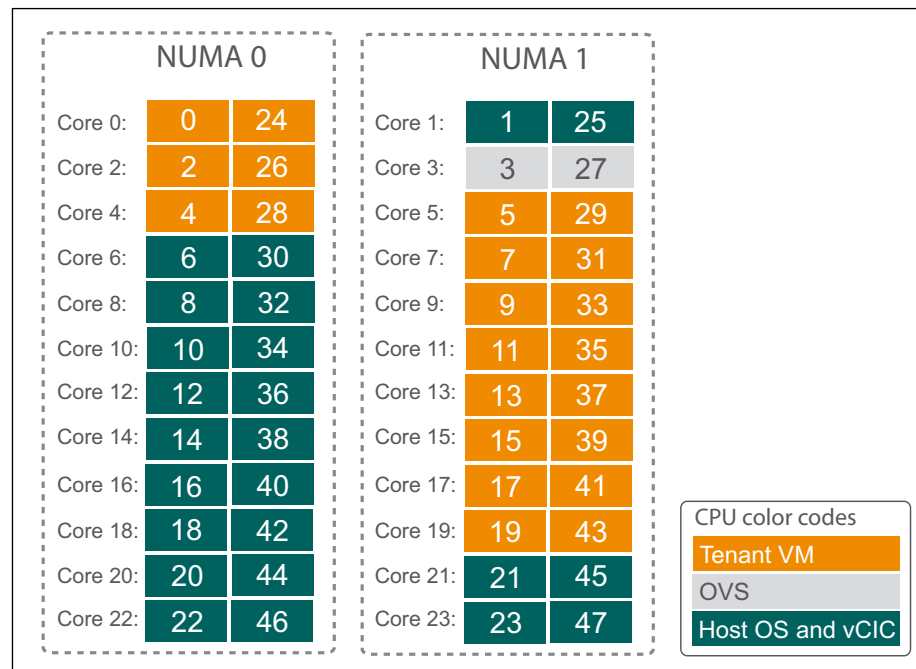
| CPU Owner                          | Allocated CPU ID  |
|------------------------------------|---|
| Tenant VM                          | 0,24, 2,26, 4,28, 5,29, 7,31, 9,33, 11,35, 13,37, 15,39, 17,41, 19,43         |
| OVS <sup>(1)(2)</sup>              | 3,27  |
| OVS control process <sup>(3)</sup> | 1 <sup>(3)</sup>  |
| vCIC <sup>(4)</sup>                | Uses non-isolated CPUs, see Host OS   |
| Host OS                            | 6,30, 8,32, 10,34, 12,36, 14,38 16,40 18,42, 20,44, 22,46, 1,25, 21,45, 23,47 |

(1) OVS configuration requires at least one PMD thread on each NUMA node with a physical interface, and at least one PMD thread on the NUMA node where the control threads are located.

(2) To achieve a more predictable performance, allocate only one CPU per core for OVS. See Section 3.3.1 on page 6 for more information.

(3) The process does not get a CPU for its exclusive use. A configuration parameter specifies one of the host OS CPUs to be used by the OVS control process. The CPU must be selected from a NUMA node to which a physical interface for the tenant network is attached.

(4) The vCIC can be allocated on cores that are over allocated and shared with the application. In such cases, it must be ensured that the application sharing resources with the vCIC don't exhaust the vCIC resources.



**Figure 2 CPU Allocation of the Respective Resource Owner per NUMA Node on Dell R630 with Intel Xeon E5-2680 v3 Processor**



### 3.3.1 Allocating Single CPUs for OVS

To achieve a more predictable OVS performance, only a single CPU must be allocated to OVS from each CPU core reserved for OVS PMD threads. The other CPU on the same core, also called hyper-thread sibling, must be isolated, that is, not used by the host OS process scheduler, and is free from any extra load that can negatively influence the OVS performance.

To achieve the described allocation, the following settings must be performed:

- Refer to the *Configuration File Guide* for the following setting in the configuration file:

The other CPU in the core must be reserved for the `idle` owner. By this reservation, the following results are achieved:

- The idle CPUs are listed in the kernel's `isolcpus` boot parameter. It ensures that the process scheduler does not assign any process to these CPUs.
- The idle CPUs are listed in the kernel's `nohz_full` boot parameter. It ensures that the kernel does not generate scheduling clock interrupts on these CPUs.

- Refer to *SW Installation in Single Server Deployment* for the following setting at the Compute host:

The idle CPUs must be listed in the `noirqs` parameter of the `configure-interrupts` upstart task in order to avoid these CPUs from using them for serving interrupt requests.

## 3.4 RAM Configuration

This section describes the optimal RAM configuration for the Single Server CEE.

Refer to the *Configuration File Guide* for more information about the configuration procedure.

The subsections contain the following information:

- Section 3.4.1 on page 6 provides general information about RAM configuration in CEE.
- Section 3.4.2 on page 7 describes the memory sizes for RAM allocation.

### 3.4.1 Introduction

The default configuration for the Host OS on the Compute node is 8 GiB. More memory can be reserved for the Host OS by setting the relevant configuration



parameters. The needed amount of memory depends on the values of a number of other parameters as described below.

The memory used for the VMs is allocated to huge pages. This is the memory visible from the inside of the VMs. The 1 GiB huge pages are referred to as **Tenant VM** in the RAM reservation tables of Section 3.4.2 on page 7. In addition to the 1 GiB huge pages, the VMs need memory allocated from the host OS. This memory is used, for example, to emulate devices used by the virtual machine. It is hard to predict the amount of host OS memory used by the emulator since, for example, it depends on the type and number of the used devices. A small VM consumes less than 100 MiB, while it can grow to several hundred MiB in specific cases. About 300 MiB host OS memory would be enough for each virtual machine but we must double it and calculate with 600 MiB as explained below.

In a system using the NUMA architecture, the NUMA location of VMs must be considered. The available memory, that is, the huge pages and the Host OS memory, are evenly distributed between the NUMA nodes. By design, OpenStack Nova allocates VMs on the first NUMA node that fits the VM. Apart from the VMs running on both NUMA nodes, the VMs allocate memory from the NUMA node on which they are running. In a worst case scenario where all VMs are allocated on the same NUMA node, all the memory for the VMs will be allocated from the same NUMA node. In such a scenario most of the memory on the other NUMA node will be unused, and half of the memory on the Compute node will be free. To be on the safe side in a dual socket system, the 300 MiB host OS memory per VM must be doubled to cover the case where all VMs are allocated on the same NUMA node.

The host OS memory usage for processes other than the VMs depends on the CPU reservations. The host OS uses the unreserved CPUs. By default, CEE allocates two cores to the host OS on NUMA node 0. This means that most of the memory used by the Host OS will be allocated from NUMA node 0. In some scenarios it is preferred to run the host OS on both NUMA nodes. It can be achieved by modifying the CPU reservation.

**Note:** In order to allocate as little memory for the host OS as possible, memory profiling of the host OS for the specific scenario is recommended.

### 3.4.2 Configuration

The minimal RAM size is 128 GiB, Table 4 specifies the volumes allocated to the resource owners.

14 GiB RAM is allocated to the host OS, so this table contains 114 GiB.

Table 5 contains the total memory sizes.

*Table 4 Memory Allocation for System with 128 GiB RAM*

| RAM Resource Owner | Hugepage size (MiB) | Number of Hugepages (Count) | Total Size (GiB) |
|--------------------|---------------------|-----------------------------|------------------|
| Tenant VM          | 1024                | 101                         | 101              |
| OVS                | 2                   | 1024                        | 2                |
| vCIC               | 1024                | 11                          | 11               |

*Table 5 Total Memory Sizes for System with 128 GiB RAM*

| RAM Resource Owner    | Total Size (GiB) |
|-----------------------|------------------|
| Tenant VM             | 101              |
| OVS                   | 2                |
| vCIC                  | 11               |
| Host Operating System | 14               |
| Total (with host OS)  | 128              |

If the system contains more than 128 GiB RAM, the extra memory goes to the VMs. For example, in case of 256 GiB RAM size, the VMs can use 229 GiB.

Each Neutron network created consumes RAM in the vCIC, and it influences the maximum number of virtual tenant networks. See Section 4.4.3 on page 15 for more information.

The default vCIC swap size is 512 MiB. A single server vCIC deployment requires a swap space of 5120 MiB for the vCIC. The swap space can be changed by setting the `vcic_swap_size` optional parameter. For configuration information, refer to the *Configuration File Guide*.

## 3.5 Storage Configuration

This section describes the local storage implementations, and disk requirements.

Refer to the *Configuration File Guide* for more information about the configuration procedure.

**Note:** Centralized storage is not supported for the Single Server CEE.

### 3.5.1 Local Storage Disk Space

This section lists requirements on disk space.

Data on the 1000 GiB SSD is listed in Table 6.



Table 6 Local Storage Disk Space

| Use   | Size  | Partition       | Note  |
|---|---|-----------------|---|
| Root partition of vCIC                            | 50 GiB  | /               | By default, the swap file is located on the root partition, and consumes disk from the allocated space for the root partition.  |
| Logs and core/crash dumps of vCIC                 | 40 GiB <ul style="list-style-type: none"> <li>• 10 GiB for logs</li> <li>• 30 GiB for core and crash dumps</li> </ul> | /var/log        | If the size of this area is increased, it increases the storage area for core and crash dumps. The 10 GiB for logs is a constant value.   |
| Database for OpenStack and Zabbix (MySQL) on vCIC | 40 GiB  | /var/lib/mysql  |   |
| Glance repository in Swift on vCIC                | 40 GiB  | /var/lib/glance | <p>The size might need to be adjusted depending on the amount and size of images stored in Glance.</p> <p>It includes temporary storage for SFTP/SCP transfers.</p> <p>In Single Server CEE the images are replicated in the same vCIC. Consequently, allocating 40 GiB for Glance allows a maximum of 20 GiB of images to be stored.</p> |
| Root partition of Compute host                    | 50 GiB  | /               |   |



| Use                                       | Size   | Partition | Note  |
|---|--|-----------|---|
| Logs and core/crash dumps of Compute host | 40 GiB <ul style="list-style-type: none"><li>• 10 GiB for logs</li><li>• 30 GiB for core and crash dumps</li></ul> | /var/log  | If the size of this area is increased, it increases the storage area for core and crash dumps. The 10 GiB for logs is a constant value. |
| <b>Sum for vCIC</b>                       | 170 GiB  |           | The sum of disk space used by the vCIC  |
| <b>Sum for Compute host</b>               | 90 GiB   |           | The sum of disk space used by the Compute host  |
| <b>Sum for vCIC and Compute host</b>      | 260 GiB  |           | The sum of disk space used by the vCIC and the Compute host   |

The remaining disk space is used as ephemeral storage for the VMs. The size of the ephemeral storage can be calculated by removing the storage area used by the vCIC and the Compute host from the total disk space. In case of the certified 1000 GiB total disk space, 740 GiB is available for ephemeral storage.

### 3.5.2 Disk Requirements for Atlas

When Virtual Machine (VM) images are loaded to Atlas as part of an `.ova` file, the image is temporarily stored in ephemeral storage in Atlas. To support loading of large images, the recommendation is to use 120 GiB for the Atlas ephemeral storage.

The local disk is used as ephemeral storage (no centralized storage). The Atlas VM occupies 120 GiB of the local disk on the compute node where it is running.

To reduce the disk allocated to Atlas, the size of the ephemeral disk can be reduced from 120 GiB to a minimum of 10 GiB.

**Note:** 30% of the ephemeral disk in Atlas is used as a temporary storage for `.ova` files. Consequently, the size of the ephemeral disk needs to be adjusted according to the size of `.ova` files to be loaded. Using a reduced disk size of 10 GiB implies that it can be impossible to load `.ova` files that contain images larger than 3 GiB.

### 3.5.3 Disk Requirements for Nova Snapshots

Nova snapshots are stored in the `/var/lib/glance` partition of the CIC node.

There are certain disk requirements for the Nova snapshots to work. Depending on the requirements and frequency on Nova snapshots, the system must be dimensioned with free disk space, according to the following guidelines:



- For a successful Nova snapshot, the amount of free space in disk partition `/var/lib/nova` in the Compute host must be **at least** double of the snapshot/VM size, since the snapshot is first extracted locally in the Compute node before it is uploaded to the Glance/Swift store.
- The disk space needed in the `/var/lib/nova` partition of the compute disk must have free space **at least** twice the size of VMs root disk. The reason is that during the extraction of the snapshot, first the delta of the VM disk will be extracted, after which the complete disk will be extracted.
- Disk partition `/var/lib/glance` in the CIC node must have free space **at least** twice of the root disk size of the VM, in order to accommodate the snapshot.

## 4 Characteristics

This section describes the system characteristics of CEE.

### 4.1 General System Limits

For the list of system limits, see Table 7.

*Table 7 General System Limits*

| Slogan                                     | Limit   |
|--|---|
| Number of hosts (servers)                  | One server is used.                               |
| Number of cores occupied by infrastructure | See Table 3 for information about CPU allocation. |

### 4.2 Orchestration Interface

The system limits for orchestration are listed in Table 8.

*Table 8 Orchestration Limits*

| Slogan            | Limits   |
|-------------------|--|
| Number of tenants | The maximum number of supported tenants is 50. |



## 4.3 Tenant Execution Environment

This section describes the tenant-related limits on the environment.

### 4.3.1 Performance

Performance limits are listed in Table 9.

*Table 9 Tenant Execution Performance*

| Slogan           | Limits  |
|------------------|---|
| Oversubscription | CPU overcommit: supported.<br>Memory overcommit: not supported<br>Disc overcommit: not supported. |

### 4.3.2 Resiliency

The Single Server CEE does not provide resiliency for tenant execution, due to the reduced hardware resources.

## 4.4 Network

This section lists the limits on the network.

Neutron with VLAN segmentation is used.

### 4.4.1 Performance

The switching performance of the CSS is measured by the packet rate (packets per second). The packet size has a very limited impact on the packet rate.

**Note:** Therefore, the forwarded amount of data (bit per second) increases if the packet size is increased.

The CSS executes a configurable number of threads running in endless loops, called PMD threads. Each PMD thread polls interfaces that are automatically assigned to it, processes the incoming packets, and puts them into a queue to be transmitted. The VM interfaces are polled by PMD threads located on the same NUMA node as their OVS control thread.

If the VM and the PMD thread polling the VM are located on different NUMA nodes, the maximum performance (packets per second) decreases since the packets must cross the NUMA border, and it increases the time for accessing the memory. A similar traffic capacity drop occurs if the interfaces are located on different NUMA nodes, since all the traffic must cross the NUMA border.





Table 10 shows the throughput for PHY to VM traffic cases, and Table 11 for VM to VM cases. Table 12 provides dimensioning guidelines for specifying the amount of capacity that is safe to use.

The assessments are based on measurements performed on a multi-blade Dell configuration.

*Table 10 Assessed per Host Forwarding Capacity, 64 Byte Frames, PHY to VM*

| Slogan   | Limits   |
|--|--|
| Bidirectional traffic from PHY on NUMA 1 to VM on NUMA 0 | One PMD core is allocated to CSS on NUMA node1. One HT is used, the other is idle. Value = 3.16 Mpps |
|  | One PMD core is allocated to CSS on NUMA node 1. Both HTs are used. Value = 3.76 Mpps                |

*Table 11 Assessed Guest VM Delivery Forwarding Capacity, 64 Byte Frames, VM to VM Intra-host Traffic*

| Slogan  | Limits  |
|---|---|
| Bidirectional traffic from VM on NUMA 0 to VM on NUMA 1 | One PMD core is allocated to CSS. One HT is used, the other is idle. Value = 2.25 Mpps. |
|   | One PMD core is allocated to CSS on NUMA node 1. Both HTs are used. Value = 3.23 Mpps.  |



Table 12 Dimensioning Capacity (Bidirectional Traffic)

| Slogan   | Limits   |
|--|--|
| Total vSwitch capacity (bidirectional traffic)         | <p>The total vSwitch capacity to be used for dimensioning is 80% of the per host forwarding value above for the number of allocated cores for OVS PMD threads on NUMA node 0.</p> <p>It is different from the value of the 'Per host Forwarding Capacity', in order to take into account external effects impacting the deterministic behavior of the virtual switch. The user can use a different value tuned for a specific system configuration (recommended if VMs on NUMA node 1 communicates with other VMs on NUMA node 1), preferably based on measurements.</p>   |
| Per interface vSwitch capacity (bidirectional traffic) | <p>The maximum dimensioning limit per interface is 80% of the per-host forwarding value for one PMD core on NUMA node 1 allocated to OVS when the HT functionality is not used. If HT is used on the cores hosting OVS PMD threads the value is 50% of the per-host forwarding value.</p> <p>When more interfaces are configured than OVS-assigned PMD threads the maximum dimensioning limit per interface is reduced by an additional factor -- the number of interfaces on a NUMA node divided by the number of PMD threads on the same NUMA node, rounding any fraction to the next higher integer. Two examples: 7 interfaces and 3 PMD threads on NUMA node 0 =&gt; <math>7/3 = 2.33</math>, round up =&gt; dividend is 3 which means that the capacity figure should be divided by 3; 9 interfaces and 3 PMD threads on NUMA node 0 =&gt; <math>9/3 = 3</math>, there is no fraction so no round up =&gt; dividend is 3 which means that the capacity figure should be divided by 3.</p> <p>It is not recommended to change the per interface limit even if measurements indicate that it is possible as the behavior is highly dependent on the automatic distribution of the interfaces over the PMD threads.</p> |



#### 4.4.2 Resiliency

Resiliency is not provided by the Single Server CEE, due to the reduced hardware resources.

#### 4.4.3 Tenant Network Limitations

Limitations of the tenant network are listed in Table 13.

*Table 13 Tenant Network Limitations*

| Slogan                              | Limits  |
|-------------------------------------|---|
| Number of virtual networks          | <p>The theoretical aggregated maximum number of virtual tenant networks per CEE region is 4050. Since each Neutron network created consumes RAM in the vCIC, this theoretical maximum cannot be reached. The default configuration of RAM for vCIC allows 128 networks. Additional memory is needed if more Neutron Networks are created.</p> <p>For rough estimations, consider that 100 Neutron networks with 1 subnet and 1 port for each cost about 2 GiB memory.</p> |
| Number of vNICs per guest VM        | The maximum number of vNICs per guest VM is 10 (+ 1 Trunk vNIC).  |
| Number of Trunk vNIC attached vLANs | The number of Trunk vNIC attached vLANs is limited to 100.  |
| Number of vNICs per server          | CSS supports up to 128 vNICs per Compute host.  |
| L2 Packet MTU                       | L2 Packet MTU size is 2140 bytes.   |

### 4.5 Storage Limitations

This section describes CEE characteristics on storage.

Only local storage is supported on Single Server CEE.

For tenants, ephemeral storage (non-persistent block storage) is supported on local disks of the compute hosts.

There is no support for distributed local storage or for any shared file system in CEE.

Swift uses the Local Storage.

Object storage through Swift is only used for CEE infrastructure.



Management of VM images is supported by the OpenStack image service.

For the Single Server CEE, only “boot from image” is supported.

If data is stored on a local disk, it is erased in case of disk failure or rollback from a failed update, meaning that the VM disappears. The application must be designed accordingly.

## 4.6 In-Service Performance

This section lists the characteristics on in-service performance.

*Table 14 In-Service Performance*

| Slogan                        | Characteristics   |
|-------------------------------|---|
| Guest execution retainability | Guest execution is not interrupted at a virtual Infrastructure management cluster restart or update.  |
| Update availability           | When the update is running, OpenStack API is unavailable for about a minute. During rollback, negative response is occasionally returned.       |
| Restart availability          | It is not possible to connect to the API during the restart. The applications are designed to handle this and will not time out during restart. |

## 5 System Limitations

This section describes the system limitations in R6.

### 5.1 OpenStack Deviations

The major deviations from the OpenStack SW are:

- Floating IP
- Object Storage
- Live Migration
- Security Groups

See relevant API descriptions for more information about limitations.



## Limitations, Listed in API Documents

See the following API documents for more information on limitations:

- *Atlas OVFT API*
- *In Service Performance Northbound API*
- *Fault Management Northbound API*
- *Performance Management Northbound API*
- *Preconfigured Key Performance Indicators*
- *OpenStack API Complete Reference*

## 5.2 SW Configurations and Options

This section describes SW configurations and options.

### 5.2.1 Number of Parallel Root Volume Operations

Nova in CEE supports about 500 parallel stop/detach root volume operations.

## 5.3 Not Supported

This unsupported function is mentioned here since it is not related to any specific configuration:

- The Dashboard does not support Internet Explorer.

## 5.4 Limitations and Workarounds

Refer to the *Limitation and Workarounds for Cloud Execution Environment (CEE)*, Reference [1].

## 5.5 Update Limitations

OpenStack API is not always available during update and rollback. When the vCICs or its hosts are rebooted, the OpenStack API is unavailable up to a few minutes. A negative response is occasionally returned. For more information, see Section 4.6 on page 16.



## Reference List

- [1] *Limitation and Workarounds for Cloud Execution Environment (CEE) AZE 102 01/4 R1A, 5/109 21-AZE 102 01/4-1*