ERICSSON

CEE Technical Description

Cloud Execution Environment

TECHNICAL PRODUCT DESCRIPTION

221 02-FGC 101 3270 Uen C

Contents

1 Introduction

This document describes the R6 release of Cloud Execution Environment (CEE), which is part of the larger Ericsson Cloud System solution. CEE is based on OpenStack® software.

This document describes the following:

- Generic OpenStack cloud concepts
- Added concepts of CEE
- Overall architecture of CEE
- Actors in CEE
- Main characteristics of CEE

2 Product Overview

CEE is an Infrastructure-as-a-Service (IaaS) solution. CEE is based on OpenStack, with additional features that expand its flexibility of use and meet the needs of telecommunication service providers.

2.1 CEE

CEE is a software product, which can execute on several hardware configurations. CEE consists of the following parts:

- Compute
- Networking
- Storage

These resources are managed via the Atlas Dashboard or via CEE OpenStack Rest API and combined they are referred to as a CEE region. An overview of the CEE region is shown in Figure 1.

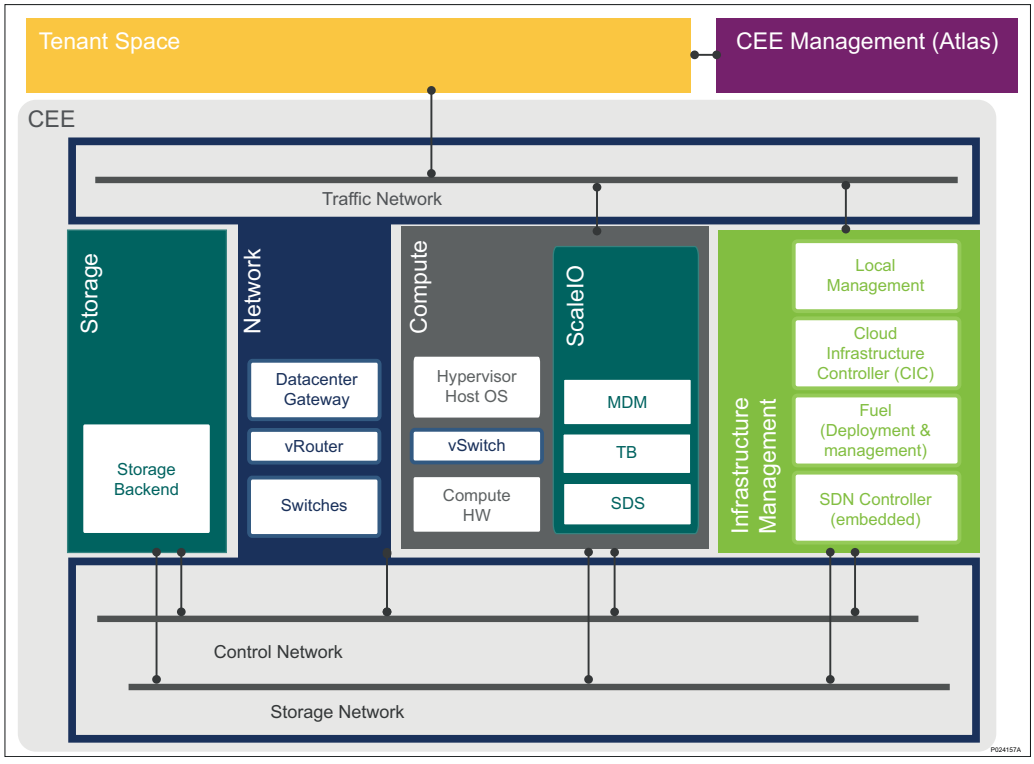

Figure 1 Overview of CEE Architecture

Ericsson has added functionality to the OpenStack components of Compute, Networking, and Storage. CEE provides the following services:

- Continuous end-user access to services, through:
  - Redundant (x3) virtual Cloud Infrastructure Controllers (vCICs)
  - Continuous Monitoring High Availability (CM-HA)
  - Persistent affinity
- High throughput with low latency, through:
  - Accelerated vSwitch (OVS)
  - Single Root Input/Output Virtualization (SR-IOV).
- Rapid and controlled deployment, through:
  - Resource aware scheduling
  - Granular resource scheduling control
  - Automated installation, with OVF support
  - Automatic and on-demand scaling
- A unified cloud infrastructure Operations and Maintenance, through:
  - Fault and performance management
  - Upgrade and rollback
- Simplified security administration
- Trusted tenant isolation

CEE is pre-verified on reference configurations. It has an efficient infrastructure utilization, and deployment options include:

- Small footprint deployments, where Virtualized controllers (vCICs) allow controllers and applications to be co-located on the same host.
- Single server, with non-redundant controller deployments

CEE contains the virtualization layer (hypervisors) and virtual switches (vSwitches). It also provides hooks to integrate and orchestrate external physical appliances, for example, physical switches and storage arrays.

Cloud management provides high-level orchestration for the following:

- Application deployment, monitoring, and management
- CEE infrastructure planning, fulfillment, assurance, and charging

Atlas and Ericsson Cloud Manager (ECM) are examples of cloud management systems. Atlas is included in CEE.

When describing the available management features in this document, the use of Atlas is assumed.

2.1.1 Hardware

The HP hardware reference configuration has the following components:

- HP c7000 system
- Extreme™ network switches
- EMC²® VNX5400 storage

The BSP 8100 hardware reference configuration has the following components:

- BSP cabinet, with six subracks

The Dell reference configuration has the following components:

- Dell R630 server system
- Extreme™ network switches
- EMC²® VNX5400 storage

The HDS reference configuration has the following components:

- Dell R630 server system
- Pluribus E28 network switches
- Juniper EX3300 control network switches

In principle, any vendor hardware can be supported, but additional integration efforts can be necessary.

2.1.2 Single Server CEE

CEE can be used in a single server deployment, using only one vCIC. There is no redundancy in compute or network infrastructure, and only local storage is supported:

- There is no High Availability for vCIC services
- Only one Compute host exists
- Ceilometer is disabled, since the statistical data is not needed

- Continuous Monitoring High Availability (CM-HA) service is disabled, so there are no alarms for CIC failed or compute host failed
- Fuel is used to install single server CEE, but it is not migrated to the single server after installation.

3 Architectural Goals and Constraints

The long-term goal for CEE is to provide support for the Infrastructure as a Service (IaaS) cloud services, which are described in this section.

3.1 Carrier Grade

The main objective of CEE is to provide carrier grade support, that is, to support the need of service providers for reliability, observability, near zero down time, security, and predictability. The characteristics of CEE services are described in the following subsections.

3.1.1 Higher Availability

CEE provides a highly available infrastructure to its tenants with nearly zero down-time.

3.1.2 Higher Security

CEE provides mechanisms which exceed security features found in traditional IT environments, since the impact is higher. In addition to standard IT security considerations, CEE addresses a large number of additional attack vectors, security policies, and runtime audits.

3.1.3 Higher Predictability

CEE is designed for real-time communication workloads. CEE provides mechanisms to ensure predictable real-time execution, low latency, low response variance, and high network throughput, even in case of small packet sizes.

3.2 Support for Distributed Cloud

Traditional IT infrastructure is focused on the Data Center. The goal is to build a scalable and flexible management infrastructure that delivers cloud-based services across multiple geographically distributed Data Centers for relatively noncomplex web and application front ends.

The Ericsson Cloud System, of which CEE is an integral part, extends the IT-grade offering with the orchestration of complex network topologies spanning across geographically distributed Data Centers. This allows flexible distribution and disaggregation of resources, such as compute servers and storage arrays.

3.3 Infrastructure as a Service

CEE supports the Infrastructure as a Service (IaaS) service model.

In IaaS, an organization outsources the equipment they use to support their own operations. Outsourced equipment includes the following:

- Hardware
- Storage
- Servers
- Network capacity

4 Functional Description

The logical components in CEE are shown in Figure 2. These are described in the following subsections.

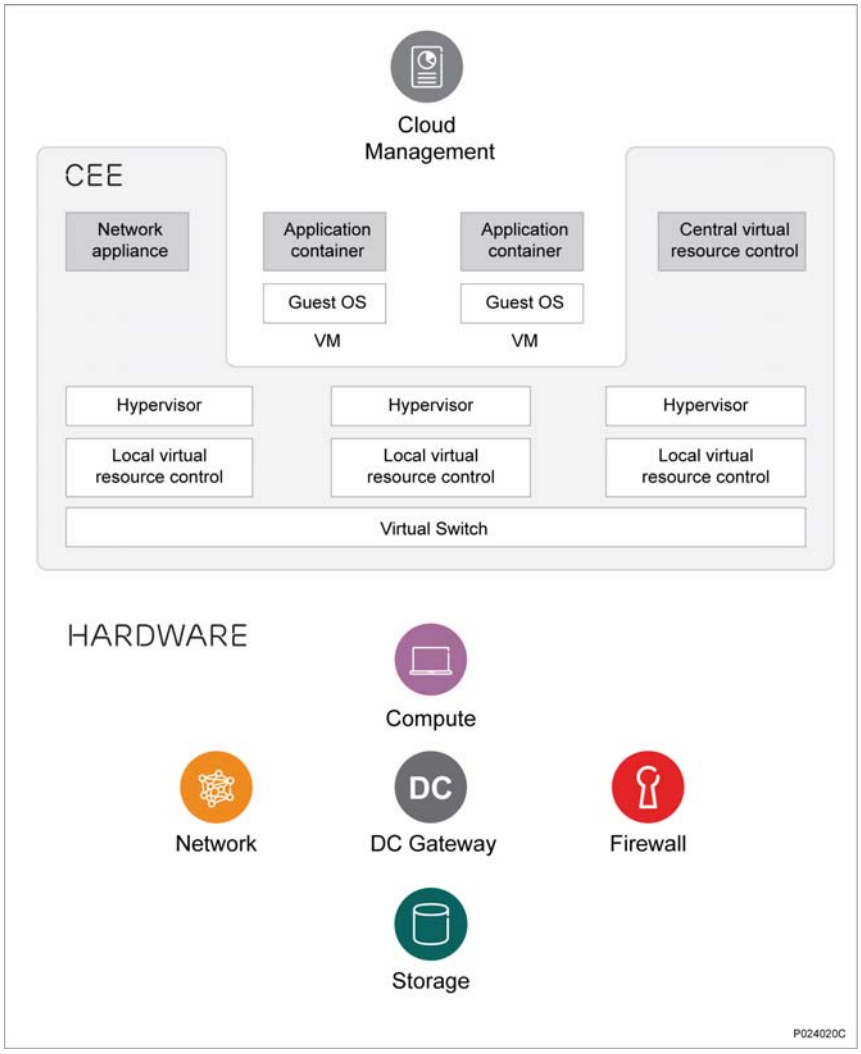

Figure 2 Logical Components in CEE

4.1 Virtual Resource Control

CEE provides the virtualization control and a management layer based on the OpenStack virtualization management system, which controls pools of compute, storage, and networking resources throughout a Data Center.

As shown in Figure 3, the following OpenStack components are used in CEE:

- Network service (Neutron) provides ‘‘network connectivity as a service’’ between interface devices managed by other OpenStack services (most likely Nova).
- Compute service (Nova) provides virtual servers upon demand.
- Image service (Glance) provides a catalog and repository for virtual disk images.
- Object Storage (Swift) allows storage and retrieval of objects (but not mounting directories like a fileserver). In CEE R6, Object Storage is not available for tenants.
- Block Storage (Cinder) provides persistent block storage to guest Virtual Machines (VMs).
- Identity management (Keystone) provides authentication and authorization for all OpenStack services.
- Performance measurement support (Ceilometer) provides counter and alarm information for charging and performance monitoring.
- Dashboard (Horizon) enables management of a CEE region. Dashboard is part of Atlas.
- Orchestration engine (Heat) launches multiple composite cloud applications based on templates. Heat is part of Atlas.
- Life cycle management (vFuel) enables management of the infrastructure hardware and software.

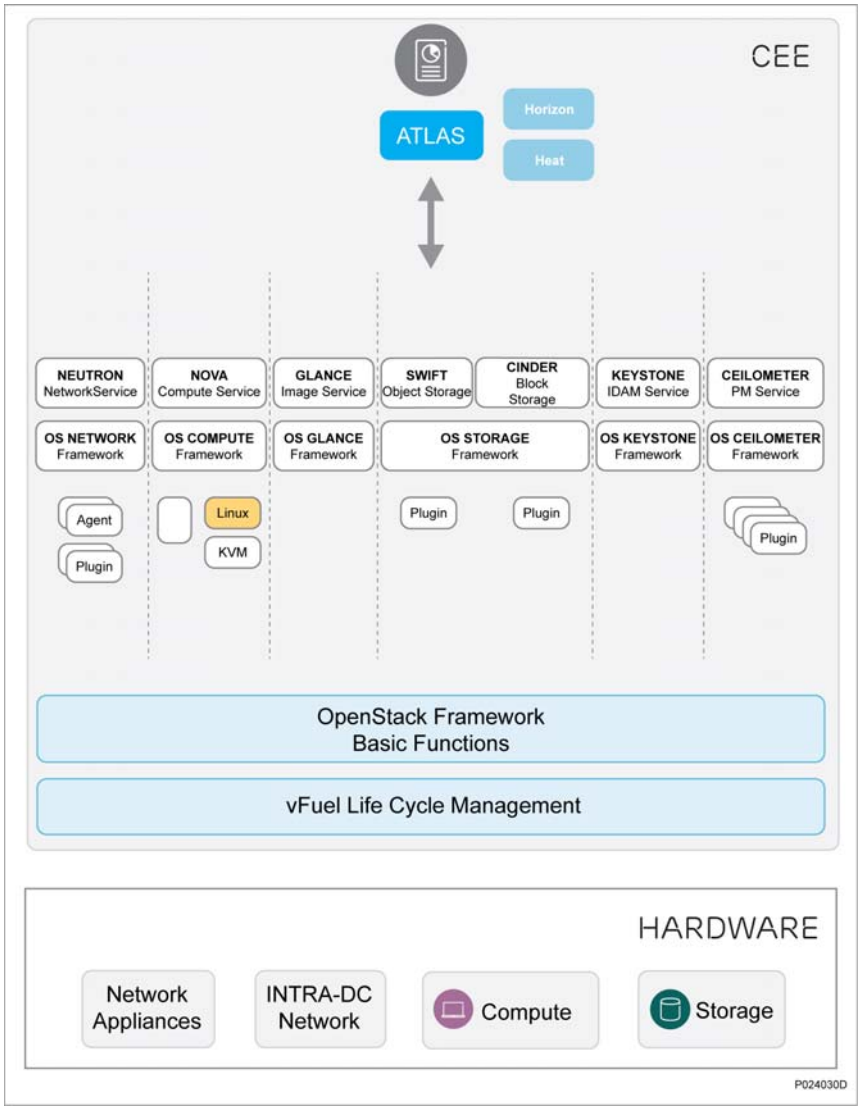

Figure 3 OpenStack in CEE

An OpenStack region is defined as one OpenStack instance running on several physical compute servers. A region is a self-contained collection of compute servers, switch fabrics, storage, and so on.

The following operations can be performed on a VM using OpenStack:

- Image Loading: when a VM is launched, a code image containing the Operating System (OS) and application that the VM is intended for is provided by the hypervisor through file system access.
- Creation of interfaces which can later be attached to a network: the VM needs to be restarted to add Virtual Network Interface Controller (vNIC) cards
- VM deployment

- Start, stop, and graceful shutdown of VMs
- Migration of VMs
- Policy-based placement
- Deployment high availability (HA) rules to achieve the required and agreed-upon level of HA for the application
- Collection of performance information from tenant VMs and infrastructure
- Reconfiguration for upgrades, failures, or optimization of system hardware use

4.2 Cloud Infrastructure Controller

The infrastructure services and control for CEE are located on three virtual Cloud Infrastructure Controller (vCIC) nodes. The vCIC nodes run on compute hosts. The three vCIC nodes form a high availability quorum.

Note: Single server deployment runs on one vCIC. There is no High Availability for vCIC services in this case.

Depending on the service, either a three-way-active service delivery model or an active-standby service delivery model applies. A summary of the internal vCIC services is shown in Figure 4.

The vCIC nodes have Ubuntu Linux as their host OS.

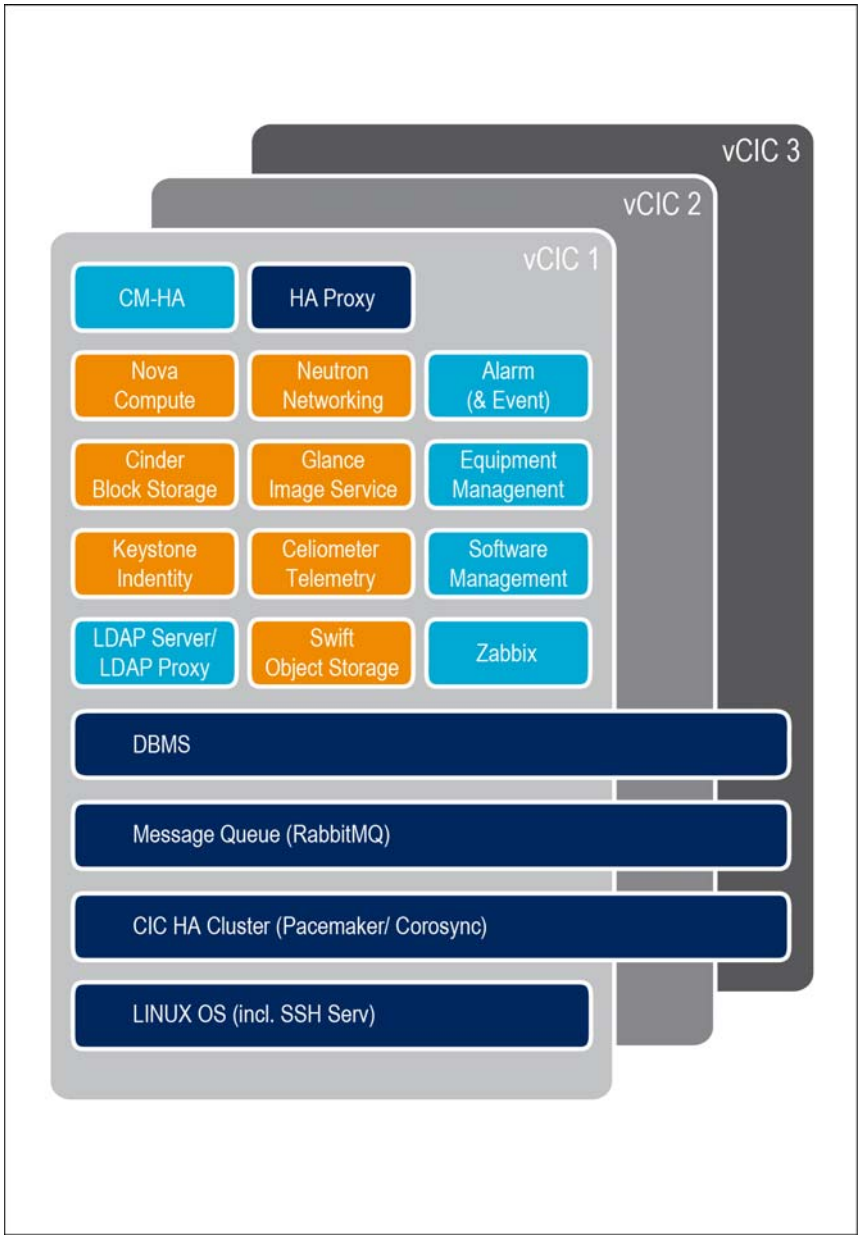

Figure 4 vCIC Redundancy Scheme

4.3 Compute Server

The compute server is connected to the switch fabric. Each compute server has a redundant connection to the switch for the VM data traffic and another redundant connection for the storage traffic.

The host OS on the compute blades is Ubuntu Linux. Linux and Windows are supported as guest OSs.

CEE uses the Kernel-Based Virtual Machine (KVM) hypervisor. KVM is optimized for the needs of telecommunication providers.

Cloud Software Defined Networking Switch (CSS) is used as a virtual switch on the compute servers.

Note: Only one Compute host exists in a single server setup

4.4 vFuel

vFuel manages the CEE infrastructure life cycle. vFuel is responsible for the following:

- Maiden software installation of the CEE software
- Updates of the CEE software
- Adding and removing hardware resources

vFuel runs as a VM on a compute host. Depending on the capacity of the compute host, vFuel and one of the vCIC nodes can run on the same compute host. vFuel uses CentOS Linux as its OS.

4.5 Multi-Server Network Configuration

The network of a multi-server CEE region is based on three switching domains, as shown in Figure 5:

Traffic switching domain
: Also referred to as Traffic Network (or traffic LAN)

Storage switching domain
: Also referred to as Storage Area Network (SAN)

Control switching domain
: Also referred to as Control Network (or control LAN)

There is physical redundancy - all physical servers are connected to all switching domains. It is fault tolerant against any Single Point of Failure (SPoF) with reasonable low failover time.

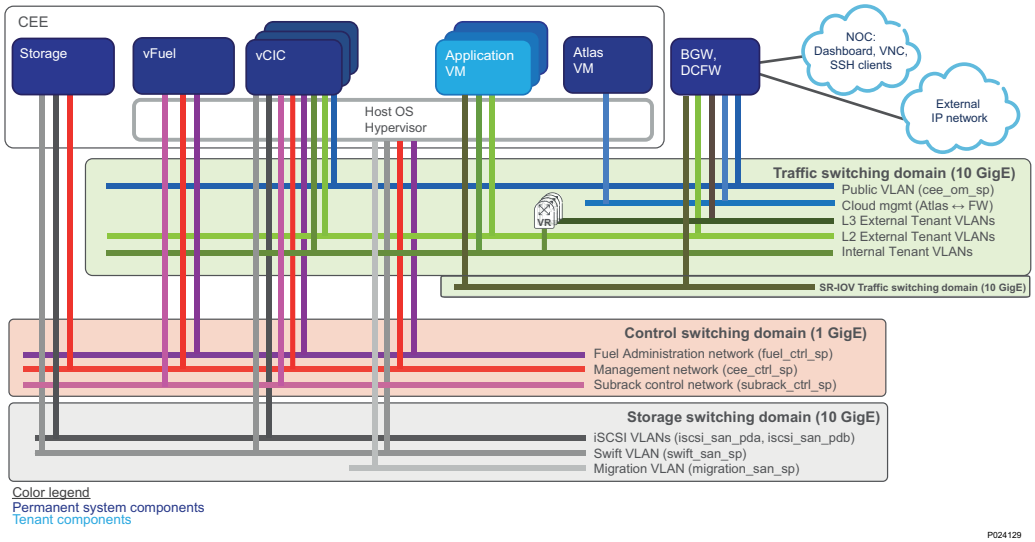

Figure 5 Logical Network Configuration, Multi-Server CEE

The **Control switching domain** is used for the OpenStack control tasks and for the management of CEE. The boot and installation of servers are also performed on this network and similarly, the upgrade of the software of the infrastructure components are also done on this network.

The most important role of the **traffic switching domain** is to forward the internal and external traffic of the tenant VMs. The traffic switching domain is the only network that has direct external connectivity. The Border Gateway (BGW) and Firewall (FW) are connected to this network. Because of the external connectivity, the northbound Operation, Administration, and Maintenance (OAM) interfaces are also connected to this network and external communication with the vCICs is also performed here.

The **storage switching domain** is used for storage access when external storage equipment is used. The external storage equipment is connected to this network using Internet Small Computer System Interface (iSCSI) protocol (`iscsi_san_pda`, `iscsi_san_pdb`). It also provides VLAN separated Migration (`migration_san_sp`) and Swift (`swift_san_sp`) related functions. The tenants do not have direct access to this network.

4.5.1 Multi-Server CEE Host Networking

From the server side, each server is connected to each switching domain with two physical Network Interface Controllers (NICs). In this document, one NIC in a pair is named *left* and the other one is named *right*. This scheme results in six physical NICs in total.In case of servers with SR-IOV traffic support, eight physical NICs are present, as shown in Figure 6. The NICs used for SR-IOV connectivity can be used from any hardware vendor in each compute server.

From the switching domain side, each domain is built from one or more network switch pairs. One *left* and one *right* physical switch form a pair, with an inter-switch link connecting the two switches.

Based on this, a server is connected to a switch pair of a switching domain by connecting the left physical NIC to the left physical switch and the right NIC to the right switch.

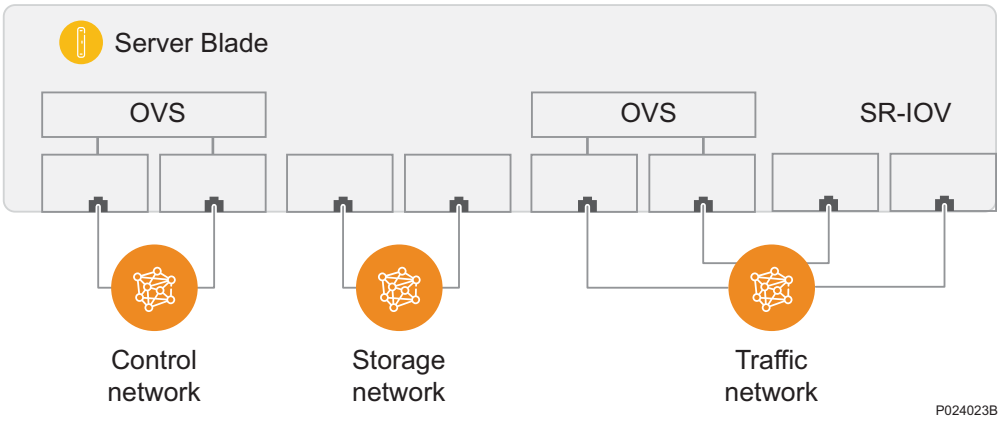

Figure 6 Aggregated Connections between Servers and Switching Domains

The two-side aggregated connectivity is introduced for reliability and performance optimization. Logically, the two sides are parts of the same domain. Consequently, the low-level virtual network architecture only comprises three switching domains with aggregated server connections.

The design of the switching domains is highly restricted by the existing physical architecture, and the resulting scheme does not provide direct connection to the logical architecture or infrastructure. For this reason, a virtual switching layer is implemented on the top of the physical NICs to provide the necessary virtual interfaces for CEE.

4.6 Single Server Network Configuration

As shown in Figure 7, the network of a single server CEE region is based on two switching domains, without redundancy:

Traffic switching domain

Also referred to as Traffic Network (or traffic LAN).

Control switching domain

Also referred to as Control Network (or control LAN).

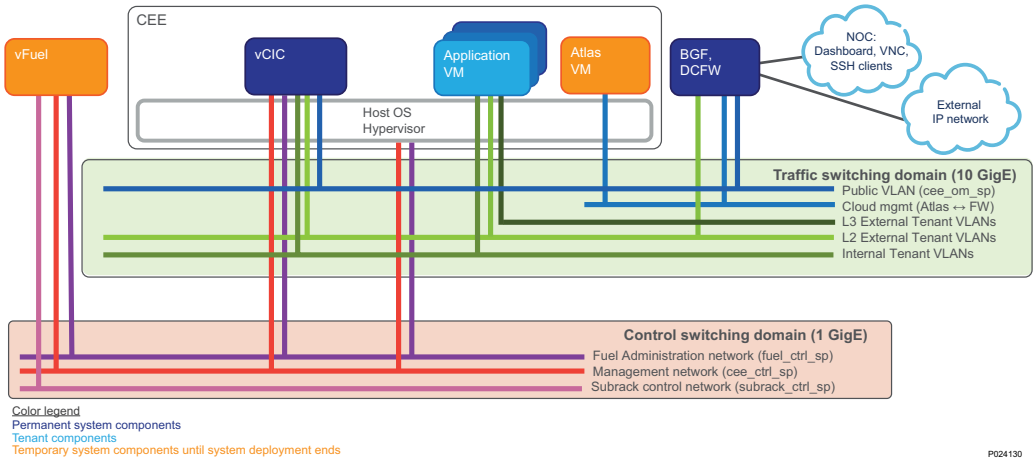

Figure 7 Logical Network Configuration, Single Server CEE

4.6.1 Single Server CEE Host Networking Solution

The two NICs are connected to the two switching domains, regardless of their role in Mirantis OpenStack.

On the control switching domain, the physical NIC is always assigned to OVS. This interface is visible to the kernel of the host operating system and it can be configured and interrogated by using traditional software tools.

4.7 Storage

CEE supports the following storage options:

Persistent storage

- OpenStack Block Storage (Cinder), see Section 4.7.1 on page 16
- OpenStack Object Storage (Swift), see Section 4.7.2 on page 17

Non-persistent storage

Ephemeral storage controlled by OpenStack Compute (Nova)

Local Storage Local, ephemeral storage.

4.7.1 OpenStack Block Storage

Some CEE configurations provide a back end for OpenStack Block Storage using one of the following solutions:

- EMC² VNX centralized storage system

- EMC² ScaleIO dedicated server solution

Note: OpenStack Block Storage is not possible without EMC² VNX centralized storage system or EMC² ScaleIO dedicated server solution. EMC² VNX and EMC² ScaleIO are mutually exclusive and optional.

OpenStack Block Storage provides persistent block level storage devices for use with OpenStack Compute instances. Block storage volumes are fully integrated into OpenStack Compute and the Dashboard, which allows cloud users to manage their own storage needs.

CEE provides a robust iSCSI multipath configuration in active/active mode between the hosts and the centralized storage system to achieve resiliency and load balancing.

4.7.2 OpenStack Object Storage

OpenStack Object Storage is supported only as a back-end for OpenStack Image Service (Glance) and CEE internal purposes. Each vCIC works as a Swift proxy and Swift storage node. Swift storage is provided by the local disks of the vCICs. Replicas are stored for each vCIC.

4.7.3 EMC ScaleIO

ScaleIO is a software-only solution that uses LAN and existing local disks of a server to create a virtual SAN that has all the benefits of external storage. ScaleIO utilizes the existing local storage devices to turn them into shared block storage.

The lightweight ScaleIO software components are installed on the application servers in a distributed server configuration. These communicate through a standard LAN to handle the application I/O requests sent to ScaleIO block volumes. The software immediately responds to the changes, rebalancing the storage distribution and achieving a layout that optimally suits the new configuration.

ScaleIO architecture in a dedicated server configuration is shown in Figure 8.

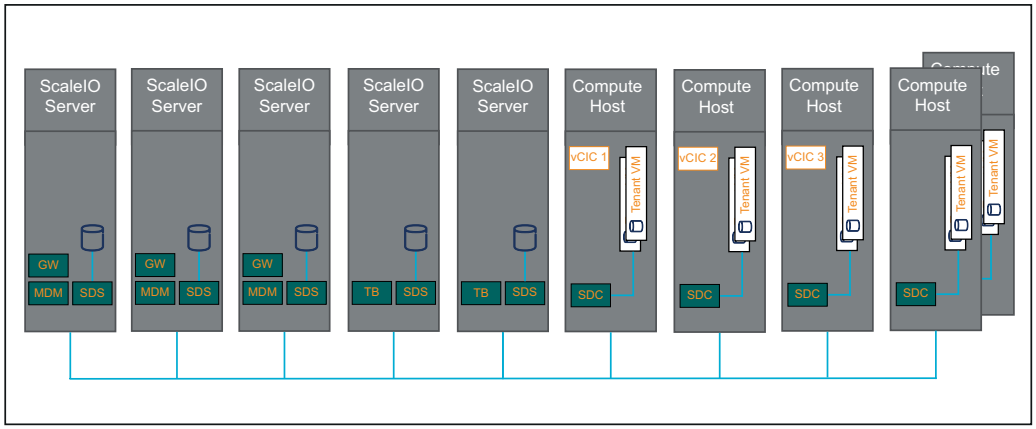

Figure 8 ScaleIO Architecture on Compute Hosts

4.8 Data Center Gateway

A Data Center gateway is required, but it is not part of CEE.

To produce a homogeneous operating environment, the Data Center gateway provides Layer 2 and Layer 3 connectivity to the Enterprise VPN and the private DataCenter, and Layer 3 connectivity to the public Internet.

Instead of Data Center Gateway, the term BGW is used in other Ericsson documents.

4.9 Data Center FW

Data Center FWs are required, but they are not part of CEE. To support configuration, CEE provides hardening guidelines for the configuration of Data Center FWs.

In the cloud infrastructure, FWs in different locations have different functionality and capacities. Examples include the following:

- Default Access Control List (ACL) with stateless rule
- ACLs with Malware Detection
- Intrusion Detection System (IDS)
- Intrusion Prevention System (IPS)

A logical view of the FW configuration for CEE is shown in Figure 9.

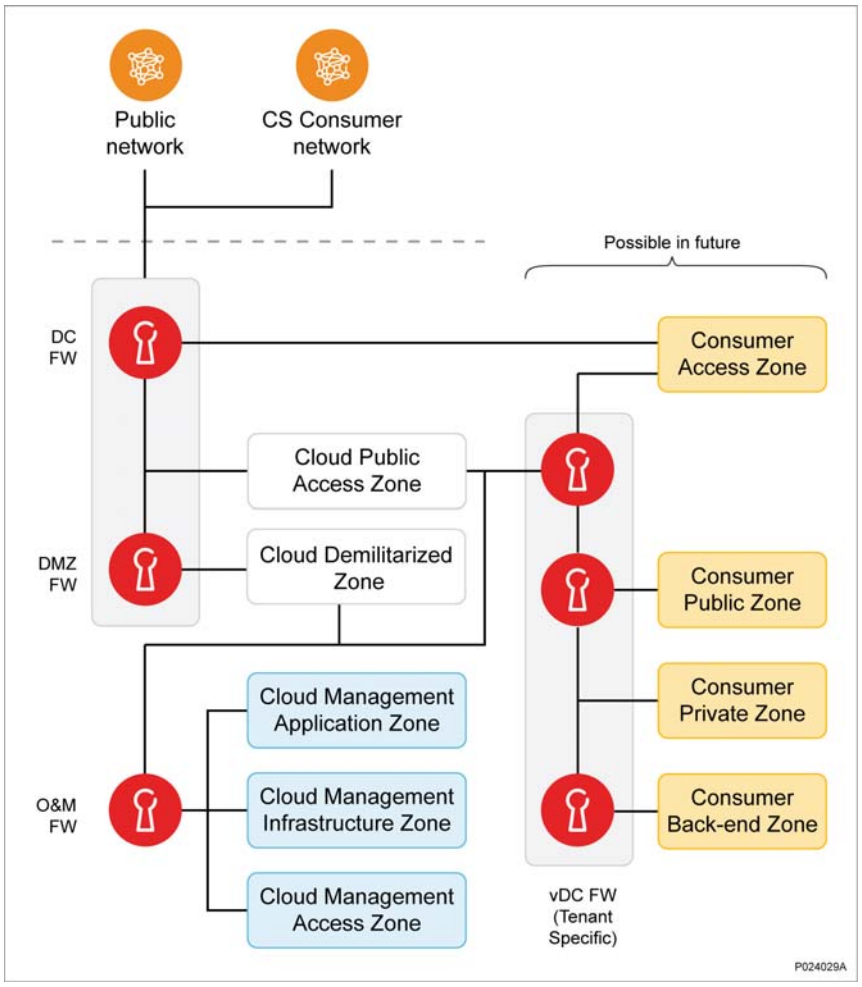

Figure 9 Logical View of FW Functionality

The logical FWs, as shown in Figure 9, are as follows:

- Data Center FW
  This FW protects the Data Center from external attacks. Communication between different Cloud Service Providers within the cloud must also go through this FW. The Data Center FW needs to support multi-Gbps traffic with Distributed Denial of Service (DDoS) protection.

- Demilitarized Zone (DMZ) FW
  This FW protects the DMZ zone.

- Cloud Operation & Maintenance (O&M) FW
  This FW protects the O&M cloud management infrastructure. The function does not require high capacity. Firewalling functionality needs to be supported, including the following:

  – Application firewalling

  - Content filtering
  - Malware detection
  - User based security policies of the cloud service

- Virtual Data Center (vDC) FW
  This FW is specific to the Cloud Service Provider (at least one FW per provider is needed). It supports the security policy of the Service Provider. The administrative control of this FW is partly delegated to the tenant administrators. There is also a back-end vDC FW, which protects the back-end network of the virtual Data Center, including the databases.

4.10 Traffic Networking

This section describes the tenant view of traffic networking.

CEE supports tenant network orchestration with OpenStack Networking (Neutron).

The performance of contemporary Data Centers greatly depends on the networking technology used to interconnect computing, storage, and services elements, which are always present in a cloud computing environment. Networking is also important because a uniform network provides access to most available resources (main processing memory yet excluded), thus East-West connectivity becomes as important as North-South. This puts high demands on the network fabric and requires non-blocking, lossless, and fine grained back-pressure characteristics.

The cloud networking solution is based on a combination of the following Layer 2 to Layer 7 elements:

- Classification
- Switching and routing
- Forwarding
- Filtering
- Shaping and policing

These elements are implemented in a combination of software and hardware.

4.10.1 Network Type 1: Internal Layer 2 Neutron Network

See Figure 10 for Network Type 1, which is an internal Layer 2 Neutron network.

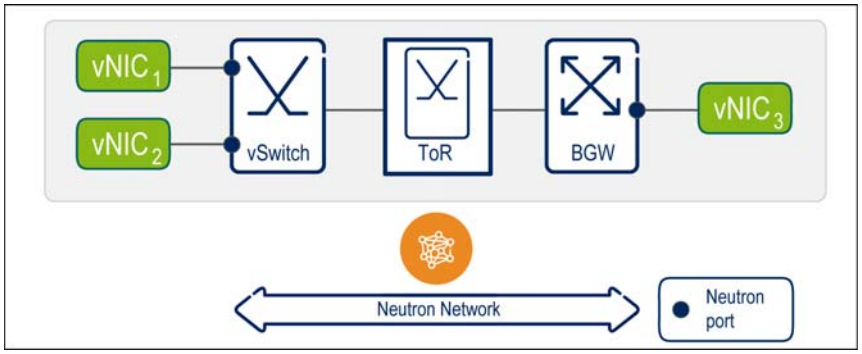

Figure 10 Network Type 1: Internal Layer 2 Neutron Network

Figure 10 shows the common setup of Layer 2 communication between two VMs in the same CEE region. Depending on whether the VMs are located on the same server or not, the Top of Rack (ToR) switch is an optional component. The two Neutron ports are defined in Neutron for each vNIC, and then associated with a common Neutron network. The Layer 2 network is realized by using a Virtual Local Area Network (VLAN).

4.10.2 Network Type 2: External Layer 2 Neutron Network

See Figure 11 for Network Type 2, which is an external Layer 2 Neutron network.

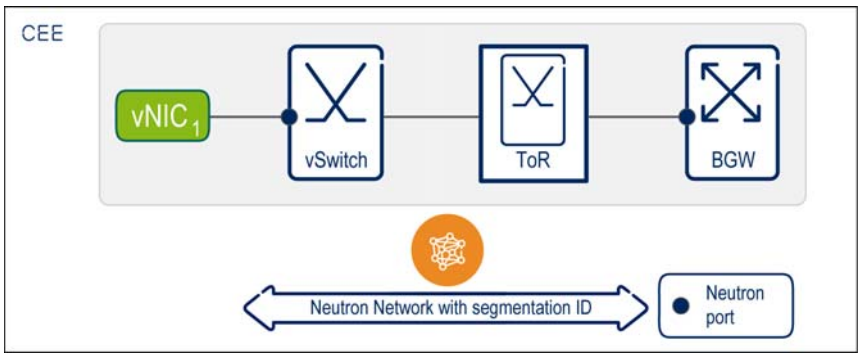

Figure 11 Network Type 2: External Layer 2 Neutron Network

Network Type 2 in Figure 11 is similar to the network in Figure 10, but Network Type 2 provides external Layer 2 connectivity towards the Data Center Gateway (BGW in the figure). There is no configuration of the actual BGW made by Neutron or CEE. In this network setup, Neutron creates Layer 2 paths in the vSwitch and ToR switch between the vNIC of the guest VM and the BGW.

4.10.3 Network Type 3: External Layer 3 Neutron Network

See Figure 12 for Network Type 3, which is an external Layer 3 Neutron network with Neutron IPv4 router and Neutron IPv4 subnets.

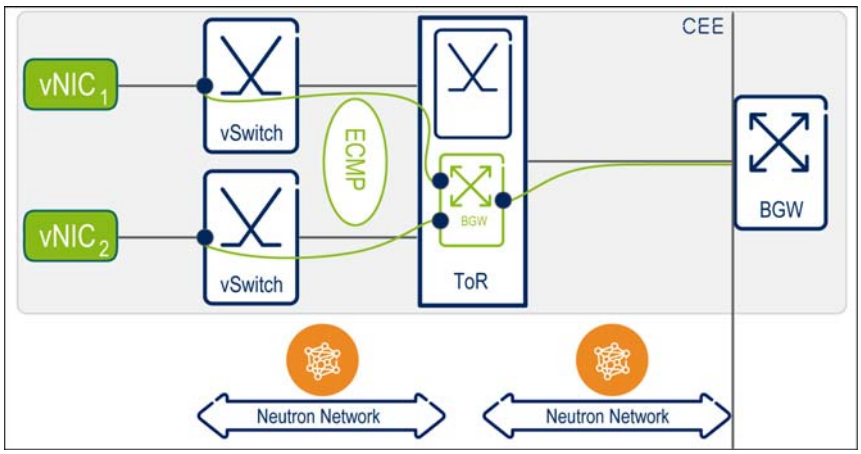

Figure 12 Network Type 3: External Layer 3 Neutron Network

Network Type 3, shown in Figure 12, provides Layer 3 connectivity for VMs. The Layer 3 routing function is carried out by virtual routers instantiated in the ToR switch. If several routes have equal cost to the same destination IP address, the virtual router in the ToR switch uses Equal-Cost Multi-Path Routing (ECMP) to distribute the Layer 3 traffic to the addressed VM. The networks are configured through Neutron, but the BGW must also be configured through its own management system (both Layer 2 and Layer 3 layers) to match the configured Neutron network. Everything configured through Neutron commands or APIs is also realized by Neutron, towards vSwitches and the ToR switch.

With the exception of the IP address assignment, Layer 3 configuration functionality is only provided in CEE configurations where the ToR switch support such actions from OpenStack Neutron.

4.10.4 Network Type 4: External Layer 2 Neutron Networks with Trunk Port

See Figure 13 for Network Type 4, external Layer 2 Neutron networks with trunk port.

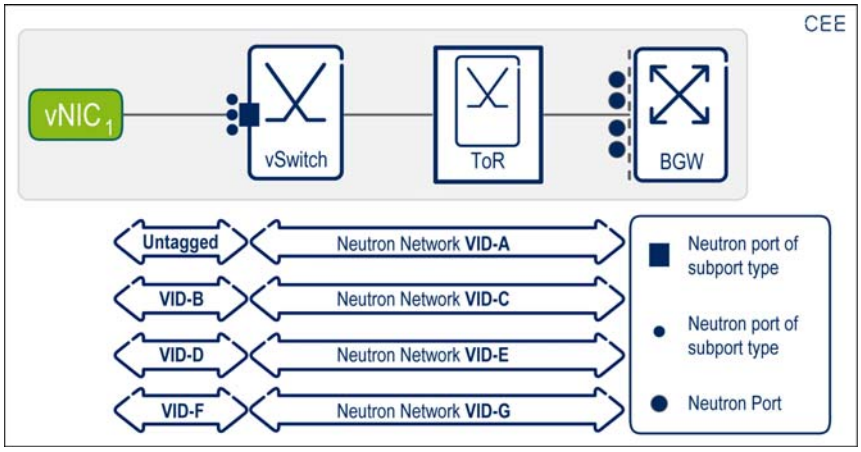

Figure 13 Network Type 4: External Layer 2 Neutron Networks with Trunkport

Network Type 4, shown in Figure 13, is a variant of the network described in Section 4.10.2 on page 21. The trunk port feature is developed by Ericsson, but currently it is not part of the OpenStack community. (The Blueprint has been approved and implementation is currently ongoing in OpenStack). The main purposes of the feature are as follows:

- To enable each VM to use several VLANs on the same Neutron (trunk) port. This makes porting of an existing legacy application into a VM environment easier.
- To enable the same VLANs to run in more than one VMs, without interference. To do that, the trunk port feature, inside the vSwitch, translates the visible guest VLANs into infrastructure-unique VLANs. This can be an advantage in the following cases:
  - If several instances of the same application run in one CEE region
  - If different applications overlap in their own respective VLAN configuration

4.10.5 Cloud Networking Model

The cloud networking model is hierarchical, with a clear relationship between the main connectivity primitives.

As shown in Figure 14, a set of network attachment points are connected by a VLAN as a virtual broadcast segment. A VLAN corresponds to a subnet at Layer 3, and one or more subnets can be combined into a Layer 3 VPN instance. The access to the public Internet is logically modeled as a VPN.

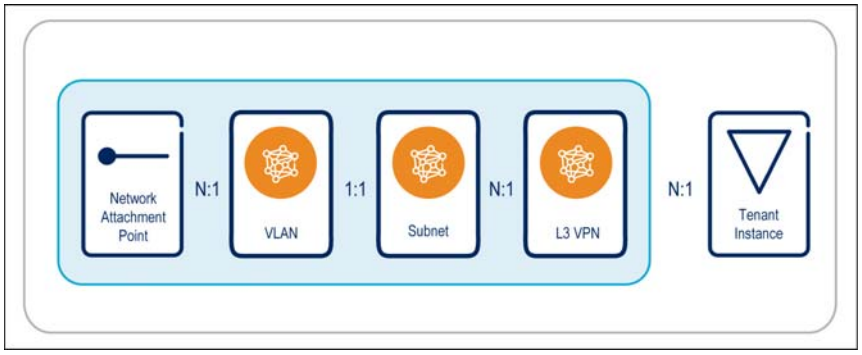

Figure 14 Cloud Networking Model

Multiple Layer 3 VPN instances constitute a Cloud Service Provider instance, although in some circumstances contiguous connectivity cannot exist across Layer 3 VPN instances. (For example, one VPN is access to the Internet, and another VPN is access to storage, and application software is connected to both of them by logically disjoint network attachment points). Connectivity can exist in the form of a FW, NAT or other filtering network appliance.

4.10.6 Layer 2 Cloud Network Connectivity

The Ethernet VLAN, which is a virtualized broadcast domain at Layer 2, and a subnet at Layer 3, is the basic unit of intra-cloud connectivity, virtualization, and Cloud Service Provider isolation.

4.10.7 Layer 3 Cloud Network Connectivity

VLAN as a subnet is the basic Layer 3 unit of connectivity, where one or more subnets can be combined to create Layer 3 VPNs. The ethertype in the Ethernet frame permits them to coexist simultaneously.

The subnets belong to three classes:

- Publically routable prefixes
- VPN (private) prefixes
- Zeroconf prefixes (requires reachability only within a cloud subnet)

A physical point of attachment to the fabric can have multiple virtual interfaces that connect to multiple classes of Layer 3 subnets. For example, a VM can have a virtual interface to a Network Address Translated (NATTed) prefix for communication with clients in the public Internet and a virtual interface to a zeroconf prefix that connects it to a set of supporting VMs.

4.10.8 Cloud Service Provider Isolation

Cloud Service Providers co-located in the same Data Center must be isolated from each other. Multiple technologies exist to support Service Provider isolation, as shown in Figure 15. Currently only one isolation technique is supported, the Customer VLAN (C-VLAN) Cloud Service Provider grouping.

The VLAN identifier (VID) space (4096) is divided among Cloud Service Provider tenants. This means that the number of private VLANs used by each Cloud Service Provider is limited to a fraction of 4096.

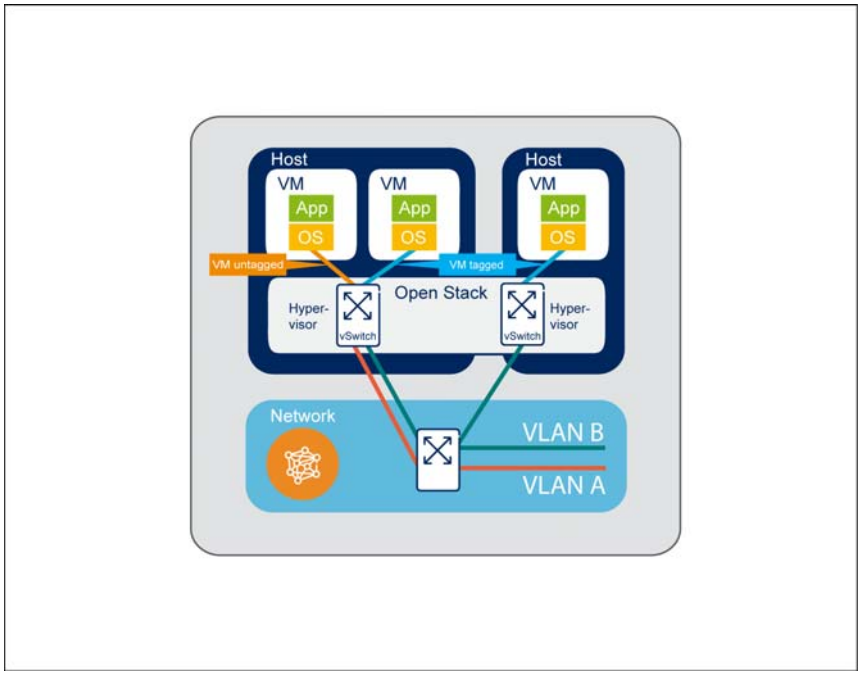

Figure 15 Cloud Service Provider Isolation

4.10.9 Resiliency

This section describes the switching resiliency in CEE.

4.10.9.1 Layer 2 Resiliency

NIC teaming is used on the two Ethernet interfaces in the traffic network domain on each host. Depending on the functionality provided by the ToR switch, different Layer 2 resiliency mechanisms of the Ethernet interfaces are used:

Link Aggregation Control Protocol (LACP) based NIC teaming:
In this configuration, network resiliency of the Ethernet interface in the servers is based on NIC teaming that uses LACP procedures. LACP is available in the host when the support in the ToR switch is based on the Extreme Multi-Chassis Link Aggregation Group (MC-LAG).

CFM based NIC teaming:
In this configuration, NIC teaming is based on Connectivity Fault Management (CFM) procedures. This is a software implementation where monitoring is performed between the CEE hosts independently of the support in the ToR switch. This option is available, for instance, when the network fabric uses stacked switches, where faults can occur between switches which are not propagated to the switch connecting to the host.

4.10.9.2 Layer 3 Resiliency

For Layer 3 resiliency, in configurations based on Extreme switches, CEE uses Virtual Router Redundancy Protocol (VRRP) for redundancy between Virtual Routers.

4.11 Cloud SDN Switch

The Cloud SDN Switch (CSS) is a network software switch based on Open Virtual Switch (OVS).

CSS consists of two main components – data plane and control plane. These components can be further divided into the following components:

- Control plane
- Command-Line Interface (CLI) tools for vSwitch configuration and troubleshooting
- Kernel data plane module for the control network
- Userspace data plane implemented using Intel® Data Plane Development Kit (DPDK) for high performance payload traffic

Control traffic for the node uses the kernel data plane. The userspace data plane gives higher capacity and is used in virtualization environments to bridge traffic between VMs and the external physical network.

The kernel data plane and the userspace dataplane are used simultaneously The CSS configuration commands indicate which datapath must be used for a particular bridge or port.

CEE uses the OpenStack framework to manage the virtualization environment: configuration of CSS is done by OpenStack network configuration. The VMs are connected to the CSS data plane by standard Virtio drivers, providing isolation between tenants.

4.12 Cloud SDN Controller

The Ericsson CSC is a Java-based SDN controller, based on the Beryllium release of the OpenDaylight (ODL) controller. The CSC provides an abstract view of the data plane that is distributed across different CSSs. The data plane processes packets across the network, subnets and L3 VPN layers. The CSC node provisions the data path using OpenFlow 1.3 interfaces on the CSSs. The CSC controls the communication across the CSSs using various overlay technologies including GRE and VXLAN.

The CSC, together with CEE, provisions the L2 and L3 VPN services. A Neutron plug-in provides the mapping to the CEE (OpenStack) Networking

API. A BGP stack is incorporated into the CSC to provide exterior gateway protocol services.

CSC utilizes HA mechanisms, fault management and performance management supported by CEE.

4.13 Switching Fabric

The CEE region reference configuration, based on Extreme switching equipment, includes:

- One pair of ToR switches for the control network
- One pair of ToR switches for the traffic network
- One pair of ToR switches for the storage network

To achieve network redundancy, the CEE certified configuration uses the following to build Link Aggregation Groups (LAGs) with active LACP between the servers and the traffic switches:

- Extreme MC-LAG on the traffic switches
- CSS bonding mode 802.3ad on the servers

CEE also supports Ericsson Blade Server Platform (BSP) hardware. The switching fabric included in BSP is described in the BSP documentation.

In addition to the certified configuration, CEE can be used with Neutron managed Layer 2 switches or SDN controllers as well as with preconfigured Layer 2 switches.

4.14 Cloud Management System

Atlas is a set of management tools for CEE. It provides a web-based user interface to CEE services and application life cycle management. Atlas is based on existing open-source OpenStack components: Horizon, Heat, and Mistral. Atlas also implements a custom component, Open Virtualization Format Translator (OVFT), to facilitate OVF-based application orchestration.

All the standard OpenStack services, such as Nova and Neutron, can be managed through Atlas. However, any CEE service can be exposed and managed in the Atlas Graphical User Interface (GUI) if its interface is integrated into the Atlas GUI. The CEE service integration into the Atlas GUI is shown in Figure 16.

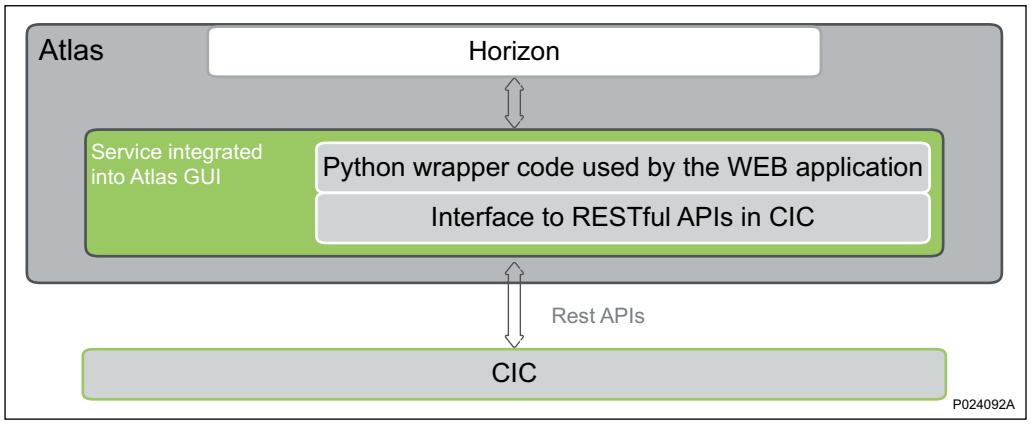

Figure 16 CEE Service Integration into Atlas GUI

The standard OpenStack Dashboard has been modified to follow the styling assets and usability guidelines described in the Ericsson User Interface Software Development Kit.

Atlas provides fault management in the form of an active alarm list. From this list, the user can get an overview of the alarm and alert history and all the active alarms in the system. The user is provided with a detailed view of the alarms, containing additional information and the Operating Instructions (OPI) of that specific alarm.

4.15 Software Management

Installation of the CEE software is done by using a kickstart server, which installs the Host OS, vCIC, and vFuel. The installation of software is done by executing commands.

The CEE software can be updated and rolled back using its configuration management software. The upgrade can be done separately for vFuel, vCIC, and compute hosts.

Note: Upgrade from 16A to R6 is not supported. CEE R6 software installation must be performed instead.

The Atlas software can be updated, upgraded, and rolled back in an existing CEE Atlas VM.

4.16 Backup and Restore

Atlas

The Atlas backup contains key configuration files. The backup is needed if the Atlas configuration must be restored to a previous state.

Fuel

Fuel is a software life cycle management application that deploys cloud environments from a single interface and enables post-deployment management of those environments. Fuel is not a highly available component of CEE. The synchronization feature makes it possible to save backups of the different stages of cloud deployment. If there is a failure, the latest state can be restored.

Note: Running Fuel inside a single server is not supported; it is not needed, since Fuel is removed from the external machine and not migrated after installation.

4.17 Audit and Health Check

A manual health check is provided to verify that CEE is running, it is available for the users, and provides the required functionality.

4.18 Performance Management

Performance data is collected by Zabbix agents from the vCIC hosts and compute hosts. The collected data is stored in a Structured Query Language (SQL) back end.

OpenStack performance data is collected by Ceilometer.

Note: Ceilometer is disabled in a single server setup, since the statistical data is not needed

4.19 High Availability

As the Ericsson cloud infrastructure is optimized for telecommunication service providers, it supports HA applications and is highly reliable. It also provides the needed APIs and components for building, deploying, and executing HA applications.

4.19.1 vCIC

To achieve high availability for the vCIC, three redundant vCIC hosts are used.

The vCIC performs a variety of services with different needs and capabilities regarding redundancy and availability. Some services run in an active-active mode, meaning that the Service Providers can be accessed through any of the vCICs. Other services run in an active-standby mode.

4.19.2 Compute

The Continuous Monitoring High Availability (CM-HA) service function adds a HA functionality for tenant VMs that is not present in a standard OpenStack environment. CM-HA uses Nova to manage VM recovery after a compute host failure.

Note: Continuous Monitoring High Availability (CM-HA) service is disabled in a single server setup

4.19.3 vFuel

vFuel is monitored by the CM-HA function.

4.19.4 Atlas

Atlas is a single VM monitored by CM-HA. Besides monitoring, CM-HA restarts, evacuates and migrates Atlas, if needed.

Internally Atlas supervises and restarts services as necessary.

4.19.5 Network

Neutron monitors and audits the Extreme switches regularly, and also makes recovery if they are wrongly configured. Recovery of a wrongly configured switch is done only if one switch differs compared to the Neutron database. The audit function is also referred to as ‘‘Consistency Checker’’.

CEE is built around three different switching domains: traffic, control, and storage. All servers are connected to all switching domains. To increase availability, all physical links and switches are duplicated.

In SAN, iSCSI multipath is used to achieve resiliency and load balancing.

4.20 Security

This section describes the security features of CEE.

4.20.1 User Management

Vanilla OpenStack Horizon user provisioning provides user administration for OpenStack administrators and tenants in the Atlas dashboard. The cloud administrator establishes tenant projects and associates users with the projects.

The purpose of CEE IdAM is to manage identities and credentials for cloud users, and to provide authentication and access control services for user accesses.

4.20.2 Security Zones

Traffic to and from CEE is passed through the Data Center gateway and the Data Center FW. All traffic identified as unwanted (based on security and network policies and rules) is dropped. The remaining traffic is handled within security zones defined by the network design.

4.20.3 Certificate Management

For secure TLS communication, the northbound interface uses a digital certificate issued by an authorized Certification Authority (CA).

4.20.4 CEE Hardening

Unnecessary services and unused ports within CEE are disabled by hardening procedures. Password policies are applied where relevant, and all user credentials can be changed.

4.20.5 Audit and Security Logging

CEE offers a logging service which records security and audit trail-related events in a central log collector inside the Atlas VM, using the Reliable Event Logging Protocol (RELP).

4.20.6 TLS on HTTPS Northbound

TLS provides confidentiality, integrity and authentication (server) between the management system and the controller for:

- OpenStack APIs
- Security and audit trail logging

4.21 End-User Access

End users can manage the virtual resources through the following interfaces:

- OpenStack northbound APIs
- GUI in Atlas
- OpenStack command line clients in Atlas
  The command line capabilities provided by the vCIC can be used by the CEE administrator for administrative tasks only.