

Release Notes

Published
2023-07-30

Contrail Networking Release 2011.L3

Table of Contents

[Introduction | 1](#)

[New and Changed Features | 1](#)

[Supported Platforms in Contrail Networking Release 2011 | 10](#)

[Container Tags | 19](#)

[Juju Charms Versions | 19](#)

[Known Behavior | 21](#)

[Resolved Issues | 51](#)

[Deprecated Issues | 51](#)

Introduction

Juniper Networks Contrail Networking is an open, standards-based software solution that delivers network virtualization and service automation for federated cloud networks. It provides self-service provisioning, improves network troubleshooting and diagnostics, and enables service chaining for dynamic application environments across enterprise virtual private cloud (VPC), managed Infrastructure as a Service (IaaS), and Networks Functions Virtualization (NFV) use cases.

These release notes accompany Release 2011.L3 of Contrail Networking. They describe new features, limitations, and known problems.

These release notes are displayed on the Contrail Networking Documentation Web page at https://www.juniper.net/documentation/en_US/contrail20/information-products/topic-collections/release-notes/index.html.

New and Changed Features

IN THIS SECTION

- [New and Changed Features in Contrail Networking Release 2011.L3 | 1](#)
- [New and Changed Features in Contrail Networking Release 2011.L2 | 2](#)
- [New and Changed Features in Contrail Networking Release 2011.L1 | 2](#)
- [New and Changed Features in Contrail Networking Release 2011 | 4](#)

The features listed in this section are new or changed as of Contrail Networking Release 2011.L3. A brief description of each new feature is included.

New and Changed Features in Contrail Networking Release 2011.L3

There are no new or changed features in Contrail Networking Release 2011.L3.

New and Changed Features in Contrail Networking Release 2011.L2

Ironic Support with OpenStack Ussuri and Ubuntu

Contrail Networking Release 2011.L2 supports OpenStack Ussuri with Ironic deployed on Ubuntu version 20.04 (Focal Fossa).

For more information, see [Ironic Support with Juju](#).

New and Changed Features in Contrail Networking Release 2011.L1

Support for Source Network Address Translation on MX Series Routers Using MS-MPC Line Cards

Starting in Contrail Networking Release 2011.L1, you can enable Source Network Address Translation (SNAT) on MX Series routers that have installed MS-MPC line cards when the MX Series router is functioning in the *DC-Gateway* fabric role. SNAT can now be used to translate the source IP addresses from physical interfaces on bare metal servers and from virtual interfaces on virtual machines for traffic traversing the MX Series router to exit a fabric.

For more information, see [Source Network Address Translation \(SNAT\)](#).

Support for Red Hat OpenShift 4.6

Starting in Contrail Networking Release 2011.L1, you can enable Contrail Networking as the Container Network Interface (CNI) in environments using Red Hat OpenShift 4.6. Red Hat OpenShift is a platform as a service product offering from Red Hat composed of containerized software elements that is orchestrated using Kubernetes. Earlier versions of Red Hat OpenShift are supported by earlier versions of Contrail Networking; see [Supported Platforms](#).

For more information on Red Hat OpenShift 4.6 in Contrail Networking, see [How to Install Contrail Networking and Red Hat OpenShift 4.6](#).

Support for Security Logging Object

Starting with Contrail Networking Release 2011.L1, you can define a security logging object (SLO) to log sessions that match a specific policy rule or security group. An SLO also enables selective session logging. This reduces the amount of data sent from vRouter agent to Contrail Analytics.

For more information, see [Security Logging Object](#).

Ironi Support with Juju

Contrail Networking Release 2011.L1 supports new charms for Ironi from OpenStack Train version 15.x.x.

For more information, see [Ironi Support with Juju](#).

Support for Configuring SNAT Logical Routers Network Data Transfer Costs

Starting with Contrail Insights Release 3.3.5, you can configure SNAT logical routers network data transfer costs. These costs are calculated as per a tiered pricing model.

For more information, see [Chargeback](#).

Support for vRouter Agent Monitoring from Contrail Insights

Contrail Insights Release 3.3.0 supports vRouter agent monitoring. A vRouter Contrail service group is also added to the Dashboard page of the Contrail Insights UI for this feature.

For more information, see [Service Monitoring from the UI](#).

Support for Ubuntu 20.04 (Focal)

Contrail Insights Release 3.3.5 supports Ubuntu 20.04 (Focal). For more information on installation, see [Contrail Insights Installation for Ubuntu Focal](#).

Octavia LBaaS Integration with Contrail Insights

Starting with Contrail Insights Release 3.3.4, Contrail Insights supports the monitoring of Octavia that provides load balancing services.

For more information, see [Service Monitoring from the UI](#).

New and Changed Features in Contrail Networking Release 2011

Collapsed Spine Architecture

Starting from release 2011, Contrail Networking supports collapsed spine, which is an architecture in which there is no defined leaf layer. In collapsed spine architecture, the Layer 3 IP-based underlay and the EVPN-VXLAN overlay, which are usually run on the leaf, are built on the spine switches.

For more information, see [Collapsed Spine Architecture](#)

Optimization of IRB Interfaces Creation in ERB Devices

In an ERB topology with **ERB-UCAST-Gateway** role configured on edge-routed bridging (ERB) switches, when creating a logical router, an integrated routing and bridging (IRB) interface is created for each associated VN on all ERB switches. IRB interfaces are created regardless of a local port or VPG in the associated VNs. This leads to a lot of unnecessary IRB and VRF configurations associated to the logical router being pushed to the ERB switches.

Starting from release 2011, Contrail Networking checks if a VPG is attached to the VN in an ERB switch with **ERB-UCAST-Gateway** role configured, before pushing IRB interface configuration associated with the logical router to the switch. Only if a VPG is attached to the VN, the IRB interface configuration is pushed to the ERB switch. The routing instance is not created if there are no VPGs.

For more information on ERB, see [Edge-Routed Bridging for QFX Series Switches](#).

Support for Fabric Telemetry Alarms

In Contrail Networking Release 2011, the Fabrics page displays additional detail about the health of your devices and interfaces, gathering data from telemetry profiles configured on these devices. These pages will show if any key performance indicators (KPIs) have crossed a threshold value. You can also configure metrics and assign telemetry profiles from the Fabrics Overview page.

For more information, see [Viewing Telemetry KPI Alarms for Fabric Devices and Ports](#).

Support for Layer 2 DCI Functionality

You can automate data center interconnect (DCI) of two different data centers. Multiple tenants connected to a logical router in a data center can exchange routes with tenants connected to a logical router in another data center. With Contrail Networking Release 2011, layer 2 DCI functionality is also

supported. You can create a layer 2 DCI of two different data centers from the **Overlay>Interconnects>Data Center Interconnect** page of the Contrail Command UI.

For more information, see [Data Center Interconnect](#).

Support for Protocol-Based Telemetry Profiles

Contrail Networking Release 2011 supports gRPC, Netconf, and SNMP protocol-based telemetry profiles. Contrail Insights collects key performance indicators (KPIs) from network devices using preconfigured sensors and commands.

For more information, see [Contrail Insights Flows in Contrail Command](#).

Support for EVPN Type-5 Routes in Logical Router

Starting from Contrail Networking Release 2011, the Contrail vRouter forwards incoming EVPN Type-5 route updates to a logical router (LR) deployed on a spine device in a fabric. This enables the LR to connect to third party physical routers and virtual machine interfaces hosted in Contrail clusters. In previous releases, the LR supports L3 traffic flow between virtual machines (VM) and bare metal servers (BMS) that are deployed in interconnected virtual networks (VN). However, when there is a default route (0.0.0.0/0) coming from a BGP peer of the tenant VN, the routing topology is disrupted due to duplicate default routes in the tenant VN VRF. In release 2011, vRouter does not add default route in VN VRF, when a virtual network is connected to a logical router.

In release 2011, the Contrail vRouter enables the LR to support inter-subnet routing between the Contrail VNs. The vRouter agent injects the non-/ 32 VN subnet route for the IPAM subnet of all the other remote VNs that are connected to the LR. This enables the vRouter agent to function without adding a 0.0.0.0/0 route to the tenant VRF.

In release 2011, Contrail Networking also provides more connectivity to customer PNF devices. An unamanged PNF device can be connected to a QFX Series device to provide connectivity to customer private networks. Instead of adding default routes, Contrail Networking adds subnet routes to the virtual networks connected to a LR. The route updates received by the logical router VRF is extended to the default EVPN route table. The Contrail vRouter agent then extends the incoming routes in all the connected tenant VRFs so that the tenant VRFs can reach the PNF subnets using a QFX Series device as next-hop.

Support for Adding or Deleting Flow Collector Nodes in Contrail Command

Contrail Networking Release 2011 supports adding, removing, and reconfiguring collector nodes (also known as *sFlow* nodes) after the system is up and running. Prior to this release, collector nodes could only be added during provisioning.

For more information, see [Adding or Deleting Flow Collector Nodes in Contrail Command](#).

Support for Intel DDP in vRouter for Fortville NICs

Starting with Contrail Networking Release 2011, the Contrail vRouter supports Intel dynamic device personalization (DDP) technology, which enables faster processing of packets with MPLSoGRE encapsulation. The Intel DDP technology is supported only in Intel Fortville Series NICs.

For more information, see [Contrail DPDK vRouter Support for Intel DDP Technology in Fortville NICs](#).

Support for Topology Discovery during Device Onboarding

Contrail Networking Release 2011 supports topology discovery when you run the onboard device job. After a device has been onboarded, you can view the fabric interfaces from the Physical Interfaces page of the device.

For more information, see [Device Import](#).

Retaining the AS Path Attribute in a Service Chain

Starting with Contrail Networking Release 2011, you can configure the AS path to be retained in the routes re-originated from the destination VN to the source VN in a service chain. You also have the ability to enable or disable the path retention for selected service chains. You can enable or disable the **Retain AS Path** option while configuring the network policy in the **Overlay > Network Policies > Create Network Policy** page.

For more information, see [Retaining the AS Path Attribute in a Service Chain](#).

Integrating OpenStack with Contrail Networking Fabric by using Modular Layer 2 (ML2) Neutron Plug-in

Contrail Networking Release 2011 supports integrating OpenStack with Contrail Networking Fabric. A Modular Layer 2 (ML2) Neutron plug-in is deployed for this integration. The ML2 plug-in enables OpenStack to use various layer 2 networking technologies that are found in complex real-world data centers.

For more information, see [Understanding OpenStack-Contrail Networking Fabric Integration](#).

Support for vRouter MAC Address - IP Address Bond Learning and BFD Health Check for Pods

Starting with Contrail Networking Release 2011, the Contrail vRouter dynamically learns the MAC address - IP address binding of the pods deployed on a virtual machine (VM). The vRouter learns the MAC address - IP address binding of the pods to enable an efficient pod to pod communication. Also,

Contrail Networking supports Bidirectional Forwarding and Detection (BFD) based health check to verify the liveliness of a pod.

For more information, see [Contrail vRouter MAC Address - IP Address Bond Learning and Bidirectional Forwarding and Detection Health Checking for Pods on Virtual Machines](#).

Support for Superspine Role

Starting with Release 2011, Contrail Networking supports the superspine role. Devices with the role superspine can function as either a lean device that provides only IP forwarding or as a route reflector that can establish iBGP sessions with the peers. To configure a device as superspine, you need to assign the Physical Role as superspine and overlay role as lean or Route-Reflector.

For more information, see [Support for Superspine Role](#).

Support for QFX5220-32CD and QFX5220-128C as Spine and Superspine Devices

Starting with Contrail Networking Release 2011, you can configure the QFX5220-32CD and QFX5220-128C devices as spine and superspine devices. They can be assigned the lean and Route-Reflector roles. The QFX5220-32CD and QFX5220-128C devices cannot be configured as leaf devices and software image upgrade is not supported on them.

For more information, see [Contrail Networking Supported Hardware Platforms and Associated Roles And Node Profiles](#) and [Support for Superspine Role](#).

Support for OpenStack Ussuri and Ubuntu Version 20.04

Starting with Contrail Networking Release 2011, Contrail Networking supports OpenStack Ussuri with Ubuntu version 18.04 (Bionic Beaver) and Ubuntu version 20.04 (Focal Fossa). The installation process is same as of installing Contrail Networking with OpenStack by Using Juju Charms.

For details, see [Installing Contrail with OpenStack by Using Juju Charms](#) and [Bundle Reference File](#).

Support for Sandump Tool on Windows Machines

Starting with Contrail Networking Release 2011, Sandump tool with Wireshark is available on Windows machines. Sandump tool captures the Sandesh messages from netlink connection between the Agent and the vRouter (only DPDK mode) and, provides detailed interpretation of all the captured bytes. For details, see [Using Sandump Tool](#).

Support for agent_header.lua Wireshark Plugin in Windows OS Computers

Starting with Contrail Networking Release 2011, you can also use the agent_header.lua Wireshark plugin in Windows OS computers, which enables you analyze the packets exchanged between vRouter data plane and vRouter agent on the pkt0 interface.

For more information, see [Using the Wireshark Plugin to Analyze Packets Between vRouter and vRouter Agent on pkt0 Interface](#).

Upgrade Contrail Networking using Red Hat Fast Forward Upgrade Procedure

Starting with Contrail Release 2011, you can use a combined procedure to upgrade Red Hat OpenStack Platform (RHOSP) from RHOSP 13 to RHOSP 16.1 by leveraging Red Hat Fast Forward Upgrade (FFU) procedure while simultaneously upgrading Contrail Networking from Release 19xx to Release 2011. For details, see [Upgrading Contrail Networking Release 19xx with RHOSP13 to Contrail Networking Release 2011 with RHOSP16.1](#).

Support for Red Hat OpenShift 4.5

Starting in Contrail Networking Release 2011, you can enable Contrail Networking as the Container Network Interface (CNI) in environments using Red Hat OpenShift 4.5. Red Hat OpenShift is a platform as a service product offering from Red Hat composed of containerized software elements that is orchestrated using Kubernetes. Earlier versions of Red Hat OpenShift are supported by earlier versions of Contrail Networking; see [Supported Platforms](#).

For more information on Red Hat OpenShift 4.5 in Contrail Networking, see [How to Install Contrail Networking and Red Hat OpenShift 4.5](#).

Support for KubeVirt in Kubernetes Environments

Starting in Contrail Networking Release 2011, you can use KubeVirt in Kubernetes-orchestrated environments that use Contrail Networking as the Container Networking Interface (CNI). KubeVirt is a virtualization add-on to Kubernetes that allows virtual machines (VMs) to run alongside the application containers present in Kubernetes environments.

For more information, see [How to Enable Virtualization with KubeVirt in Environments Using Kubernetes with a Contrail Cluster](#).

Support for Keystone Authentication in Kubernetes Environments Using Juju

Starting in Contrail Networking Release 2011, you can use the Keystone authentication service in OpenStack for authentication in environments that contain cloud networks using both Openstack and Kubernetes orchestrators when the Kubernetes environment is running Juju. This capability simplifies

authentication in mixed cloud environments and is available when the cloud networks are both using Contrail Networking.

For more information, see [How to Enable Keystone Authentication in a Juju Cluster within a Kubernetes Environment](#).

Support for Installing Contrail Insights and Contrail Insights Flows on Juju Cluster

Contrail Networking Release 2011 supports installing Contrail Insights and Contrail Insights Flows on Juju cluster after Contrail Networking and Contrail Command are installed. For more information see [How to Deploy Contrail Command and Import a Cluster Using Juju](#).

Support for Custom Names for Privileged Ports and VXLANs

In Contrail Networking Release 2011, this is a beta feature that allows you to assign custom names to privileged ports and VXLANs in order to make them easier to identify in the topology mapping. When you make queries for traffic flows in the Topology View, the custom name displays versus the port number or VXLAN ID.

For more information, see [Assign Custom Names to Privileged Ports and VXLAN IDs](#).

Support for contrail-vrouter-utils package in the Contrail Tools Container

Starting with Contrail Networking Release 2011, the **contrail-vrouter-utils** package is available only in the **contrail-tools** container. You must use the **contrail-tools** container to execute tools like **vif**, **nh**, **rt**, and so on available in the **contrail-vrouter-utils** package. In previous releases, the **contrail-vrouter-utils** package is available in the **contrail-vrouter-agent** and **contrail-vrouter-dpdk** container. You can no longer use the **contrail-vrouter-agent** and **contrail-vrouter-dpdk** containers to execute the tools available in the **contrail-vrouter-utils** package.

Support for Netronome SmartNIC vRouter for Juju Charms Deployment—Evaluation Purposes Only

Starting with Contrail Networking Release 2011, Contrail Networking supports Netronome Agilio CX for Contrail Networking deployment with Juju charms. This feature enables increased packets per second (PPS) capacity of Contrail vRouter datapath allowing applications to reach their full processing capacity. Additionally, it allows to reclaim CPU cores from Contrail vRouter off-loading permitting more VMs and VNFs to be deployed per server. For details, see [Using Netronome SmartNIC vRouter with Contrail Networking and Juju Charms](#).

Supported Platforms in Contrail Networking Release 2011

Table 1 on page 10 lists the orchestrator releases and the corresponding operating systems and kernel versions supported by Contrail Networking Release 2011.

Table 1: Supported Platforms

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
Contrail Networking Release 2011.L3	Kubernetes 1.18.17	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-154-generic MaaS Version: 2.6.2
	Kubernetes 1.14.8	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.25.1.el7.x86_64 Ansible version : 2.7.11 Docker version: 18.03.1-ce
	OpenShift 4.6.23	Operator Framework	RHEL CoreOS 4.6.8—Linux Kernel Version 4.18.0-193.47.1
	Openstack Rocky	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.25.1.el7.x86_64 Ansible version: 2.5.2 Docker version: 18.03.1-ce
	Openstack Queens	Ansible	CentOS 7.9—Linux Kernel Version 3.10.0-1160.25.1.el7.x86_64 Ansible version: 2.5.2 Docker version: 18.03.1-ce

Table 1: Supported Platforms (Continued)

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	Canonical OpenStack Train	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-154-generic MaaS Version: 2.4.2
	Canonical OpenStack Ussuri	Juju Charms	Ubuntu 20.04.3—Linux Kernel Version 5.4.0-81-generic MaaS Version: 2.4.2
	Red Hat OpenStack Platform 13.0.12/13.0.14/13.0.15	RHOSP 13 director	<ul style="list-style-type: none"> RHEL 7.8—Linux Kernel Version 3.10.0-1127 Overcloud Image: overcloud-full-13.0-20200610.2.el7ost.x86_64.tar Red Hat Content Sync Date : 2020-07-07 RHEL 7.9—Linux Kernel Version 3.10.0-1160 Director Image: octavia-amphora-image-13.0-20200930.1.el7ost.noarch Red Hat Content Sync Date : 2021-02-09 RHEL 7.9—Linux Kernel Version 3.10.0-1160.15.2 Director Image: octavia-amphora-image-13.0-20201112.1.el7ost.noarch Red Hat Content Sync Date : 2021-06-03 RHEL 7.9—Linux Kernel Version 3.10.0-1160.25.1 Director Image: octavia-amphora-image-13.0-20210511.1.el7ost.noarch Red Hat Content Sync Date : 2021-08-29

Table 1: Supported Platforms (Continued)

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	Red Hat OpenStack Platform 16.1.3/16.1.6	RHOSP 16 Director	<ul style="list-style-type: none"> RHEL 8.2—Linux Kernel Version 4.18.0-193.29.1 Director Image: rhosp-director-images-16.1-20201202.1.el8ost.noarch Red Hat Content Sync Date : 2021-02-09 RHEL 8.2—Linux Kernel Version 4.18.0-193.51.1.el8_2.x86_64 Director Image: rhosp-director-images-ipa-16.1-20210430.3.el8ost.noarch Red Hat Content Sync Date : 2021-08-18 (02:19 PM PDT)
Contrail Networking Release 2011.L2	Kubernetes 1.18.17	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-143-generic MaaS Version: 2.6.2
	Kubernetes 1.14.8	Ansible	<ul style="list-style-type: none"> CentOS 7.9—Linux Kernel Version 3.10.0-1160.25.1.el7.x86_64 Ansible version: 2.7.11 Docker version: 18.03.1-ce
	OpenShift 4.6.23	Operator Framework	RHEL CoreOS 4.6.8—Linux Kernel Version 4.18.0-193.47.1
	Openstack Rocky	Ansible	<ul style="list-style-type: none"> CentOS 7.9—Linux Kernel Version 3.10.0-1160.25.1.el7.x86_64 Ansible version: 2.5.2 Docker version: 18.03.1-ce

Table 1: Supported Platforms *(Continued)*

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	Openstack Queens	Ansible	<ul style="list-style-type: none"> CentOS 7.9—Linux Kernel Version 3.10.0-1160.25.1.el7.x86_64 <p>Ansible version: 2.5.2</p> <p>Docker version: 18.03.1-ce</p>
	Canonical OpenStack Train	Juju Charms	<ul style="list-style-type: none"> Ubuntu 18.04.5—Linux Kernel Version 4.15.0-143-generic <p>MaaS Version: 2.4.2</p>
	Canonical OpenStack Ussuri	Juju Charms	<ul style="list-style-type: none"> Ubuntu 20.04.1—Linux Kernel Version 5.4.0-77-generic <p>MaaS Version: 2.4.2</p>

Table 1: Supported Platforms (Continued)

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	Red Hat OpenStack Platform 13.0.12/13.0.14/13.0.15	RHOSP 13 director	<ul style="list-style-type: none"> • RHEL 7.8—Linux Kernel Version 3.10.0-1127 Overcloud Image: overcloud-full-13.0-20200610.2.el7ost.x86_64.tar Red Hat Content Sync Date: 2020-07-07 • RHEL 7.9—Linux Kernel Version 3.10.0-1160 Director Image: octavia-amphora-image-13.0-20200930.1.el7ost.noarch Red Hat Content Sync Date : 2021-02-09 • RHEL 7.9—Linux Kernel Version 3.10.0-1160.25.1 Director Image: octavia-amphora-image-13.0-20201112.1.el7ost.noarch Red Hat Content Sync Date : 2021-06-03
	Red Hat OpenStack Platform 16.1.3/16.1.6	RHOSP 16 Director	<ul style="list-style-type: none"> • RHEL 8.2—Linux Kernel Version 4.18.0-193.29.1 Director Image: rhosp-director-images-16.1-20201202.1.el8ost.noarch Red Hat Content Sync Date : 2021-02-09 • RHEL 8.2—Linux Kernel Version 4.18.0-193.51.1.el8_2.x86_64 Director Image: rhosp-director-images-x86_64-16.1-20210430.3.el8ost.noarch Red Hat Content Sync Date : 2021-06-10

Table 1: Supported Platforms *(Continued)*

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
Contrail Networking Release 2011.L1	Kubernetes 1.18.17	Juju Charms	Ubuntu 18.04.5—Linux Kernel Version 4.15.0-140-generic MaaS Version: 2.6.2
	Kubernetes 1.14.8	Ansible	<ul style="list-style-type: none"> CentOS 7.8—Linux Kernel Version 3.10.0-1127 Ansible version: 2.7.11 Docker version: 18.03.1-ce
	OpenShift 3.11	Ansible	RHEL 7.9—Linux Kernel Version 3.10.0-1160
	OpenShift 4.6.23	Operator Framework	RHEL CoreOS 4.6.8—Linux Kernel Version 4.18.0-193.47.1
	Openstack Rocky	Ansible	<ul style="list-style-type: none"> CentOS 7.9—Linux Kernel Version 3.10.0-1160 Ansible version: 2.5.2 Docker version: 18.03.1-ce
	Openstack Queens	Ansible	<ul style="list-style-type: none"> CentOS 7.9—Linux Kernel Version 3.10.0-1160 Ansible version: 2.5.2 Docker version: 18.03.1-ce
	Canonical OpenStack Train	Juju Charms	<ul style="list-style-type: none"> Ubuntu 18.04.5—Linux Kernel Version 4.15.0-140-generic MaaS Version: 2.4.2

Table 1: Supported Platforms *(Continued)*

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	Canonical OpenStack Ussuri	Juju Charms	<ul style="list-style-type: none"> Ubuntu 20.04.1—Linux Kernel Version 5.4.0-65-generic <p>MaaS Version: 2.4.2</p>
	Red Hat OpenStack Platform 13.0.12/13.0.14	RHOSP 13 director	<ul style="list-style-type: none"> RHEL 7.8—Linux Kernel Version 3.10.0-1127 <p>Overcloud Image: overcloud-full-13.0-20200610.2.el7ost.x86_64.tar</p> <p>Red Hat Content Sync Date: 2020-07-07</p> <ul style="list-style-type: none"> RHEL 7.9—Linux Kernel Version 3.10.0-1160 <p>Director Image: octavia-amphora-image-13.0-20200930.1.el7ost.noarch</p> <p>Red Hat Content Sync Date : 2021-02-09</p>
	Red Hat OpenStack Platform 16.1.3	RHOSP 16 Director	<ul style="list-style-type: none"> RHEL 8.2—Linux Kernel Version 4.18.0-193.29.1 <p>Director Image: rhosp-director-images-16.1-20201202.1.el8ost.noarch</p> <p>Red Hat Content Sync Date : 2021-02-09</p>
	VMware vCenter 6.7	Ansible	<ul style="list-style-type: none"> ESX version 6.5 <p>CentOS VM version running vRouter: CentOS 7.8—Linux Kernel Version 3.10.0-1160</p>

Table 1: Supported Platforms (Continued)

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
Contrail Networking Release 2011	Kubernetes 1.18	Juju Charms	<ul style="list-style-type: none"> Ubuntu 18.04.5—Linux Kernel Version 4.15.0-128-generic MaaS Version: 2.4.6
	Kubernetes 1.14.8	Ansible	<ul style="list-style-type: none"> CentOS 7.8—Linux Kernel Version 3.10.0-1160 Ansible version : 2.7.11 Docker version: 18.03.1-ce
	OpenShift 3.11	Ansible	<ul style="list-style-type: none"> RHEL 7.9—Linux Kernel Version 3.10.0-1160
	OpenShift 4.5	Operator Framework	<ul style="list-style-type: none"> RHEL CoreOS 4.5—Linux Kernel Version 4.18.0-193.13.2
	OpenStack Rocky	Ansible	<ul style="list-style-type: none"> CentOS 7.9—Linux Kernel Version 3.10.0-1160 Ansible version: 2.5.2 Docker version: 18.03.1-ce
	OpenStack Queens	Canonical OpenStack Queens	<ul style="list-style-type: none"> Ubuntu 18.04.5—Linux Kernel Version 4.15.0-126-generic MaaS Version: 2.4.2
		Juju Charms	<ul style="list-style-type: none"> Ubuntu 18.04.5—Linux Kernel Version 4.15.0-126-generic MaaS Version: 2.4.2

Table 1: Supported Platforms *(Continued)*

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	Canonical OpenStack Train	Juju Charms	<ul style="list-style-type: none"> Ubuntu 18.04.5—Linux Kernel Version 4.15.0-126-generic <p>MaaS Version: 2.4.2</p>
	Canonical OpenStack Ussuri	Juju Charms	<ul style="list-style-type: none"> Ubuntu 20.04.1—Linux Kernel Version 5.4.0-54-generic <p>MaaS Version: 2.4.2</p>
	Red Hat OpenStack Platform 13.0.12	RHOSP 13 director	<ul style="list-style-type: none"> RHEL 7.8—Linux Kernel Version 3.10.0-1127 <p>Overcloud Image: overcloud-full-13.0-20200610.2.el7ost.x86_64.tar</p> <p>Red Hat Content Sync Date : 2020-07-07</p> <ul style="list-style-type: none"> RHEL 7.9—Linux Kernel Version 3.10.0-1160 <p>Overcloud Image: overcloud-full-13.0-20200930.1.el7ost.x86_64.tar</p> <p>Red Hat Content Sync Date : 2020-10-07</p>
	Red Hat OpenStack Platform 16.1.2	RHOSP 16 Director	<ul style="list-style-type: none"> RHEL 8.2—Linux Kernel Version 4.18.0-193.19.1 <p>Overcloud Image: overcloud-full-16.1-20201003.1.el8ost.x86_64.tar</p> <p>Red Hat Content Sync Date : 2020-10-26</p>

Table 1: Supported Platforms *(Continued)*

Contrail Networking Release	Orchestrator Release	Deployment Tool	Operating System, Kernel, and Key Components Version
	VMware vCenter 6.7	Ansible	<ul style="list-style-type: none"> ESX version 6.5 <p>CentOS VM version running vRouter: CentOS 7.8—Linux Kernel Version 3.10.0-1127</p>

Container Tags

Container tags are needed to identify the files to download from the Contrail Container Registry during a Contrail Networking installation or upgrade.

To obtain the container tag for any Contrail Networking Release 20-based software, see [README Access to Contrail Registry 20XX](#).

Juju Charms Versions

The published versions of charms are available in the [Charms Store](#).

[Table 2 on page 20](#) lists the charms versions supported by the Contrail Networking Release 2011.

Table 2: Supported Charms Versions

Contrail Networking Release	Charms Version
Contrail Networking Release 2011.L3	url: cs:~juniper-os-software/contrail-agent-28 url: cs:~juniper-os-software/contrail-analytics-26 url: cs:~juniper-os-software/contrail-analyticsdb-26 url: cs:~juniper-os-software/contrail-controller-28 url: cs:~juniper-os-software/contrail-keystone-auth-27 url: cs:~juniper-os-software/contrail-openstack-29 url: cs:~juniper-os-software/contrail-kubernetes-master-22 url: cs:~juniper-os-software/contrail-kubernetes-node-20 url: cs:~juniper-os-software/contrail-command-9
Contrail Networking Release 2011.L2	url: cs:~juniper-os-software/contrail-agent-27 url: cs:~juniper-os-software/contrail-analytics-25 url: cs:~juniper-os-software/contrail-analyticsdb-25 url: cs:~juniper-os-software/contrail-controller-27 url: cs:~juniper-os-software/contrail-keystone-auth-26 url: cs:~juniper-os-software/contrail-openstack-28 url: cs:~juniper-os-software/contrail-kubernetes-master-21 url: cs:~juniper-os-software/contrail-kubernetes-node-19 url: cs:~juniper-os-software/contrail-command-8

Known Behavior

IN THIS SECTION

- [Known Behavior in Contrail Networking Release 2011.L3 | 21](#)
- [Known Behavior in Contrail Networking Release 2011.L2 | 28](#)
- [Known Behavior in Contrail Networking Release 2011.L1 | 35](#)
- [Known Behavior in Contrail Networking Release 2011 | 44](#)

This section lists known limitations with this release.

Known Behavior in Contrail Networking Release 2011.L3

OpenStack

- CEM-23142 Few deployment instances of RHOSP13 deployment with Contrail Networking 2011.L2 failed during installation of Openstack components. RHOSP13 customers are advised to contact Redhat support referring Case ID - 02995585 prior to installation.
- CEM-22783 Upgrading a RHOSP13 cluster from RHOSP13 images tested with Contrail Networking 2011.L1 to the RHOSP13 images tested with Contrail Networking 2011.L2 fails.

As a workaround, disable the Ceilometer service before upgrading. It is recommended to contact Redhat support and refer the case ID 02989025 before proceeding with upgrade.

- CEM-20280 In a Kubernetes and OpenStack joint setup, vrouter-agent restart sometimes leads to an unauthorized operation error. To resolve the issue, restart the vrouter-agent again.
- CEM-17991 In an OpenStack HA setup provisioned using Kolla and OpenStack Rocky, if you shut down all the servers at the same time and bring them up later, the Galera cluster fails. To recover the Galera cluster, follow these steps:

1. Edit the `/etc/kolla/mariadb/galera.cnf` file to remove the `wsrep` address on one of the controllers as shown here.

```
wsrep_cluster_address = gcomm://
#wsrep_cluster_address = gcomm://10.x.x.8:4567,10.x.x.10:4567,10.x.x.11:4567
```

NOTE: If all the controllers are shut down in the managed scenario at the same time, you must select the controller that was shut down last.

2. Docker start mariadb on the controller on which you edited the file.
 3. Wait for a couple of minutes, ensure that the mariadb container is not restarting, and then Docker start mariadb on the remaining controllers.
 4. Restore the `/etc/kolla/mariadb/galera.cnf` file changes and restart the mariadb container on the previously selected controller.
- CEM-15764 In Octavia Load Balancer, traffic destined to the Floating IP of the load balancer VM does not get directed to the backend VMs. Traffic destined to the actual VM IP of the Load Balancer VM will work fine.
 - CEM-9979 During upgrade of DPDK computes deployed with OOO Heat Templates in RHOSP environment, vRouter coredumps are observed. This is due to the sequence in which the services are started during upgrade and does not have impact on cluster operation.
 - CEM-5141 For deleting compute nodes, the UI workflow will not work. Instead, update the `instances.yaml` with `"ENABLE_DESTROY: True"` and `"roles:"` (leave it empty) and run the following playbooks.

```
ansible-playbook -i inventory/ -e orchestrator=openstack --tags nova playbooks/
install_openstack.yml
ansible-playbook -i inventory/ -e orchestrator=openstack playbooks/install_contrail.yml
```

For example:

```
global_configuration:
    ENABLE_DESTROY: True
...
...
instances:
```



```

...
...
  srvr5:
    provider: bms
    ip: 19x.xxx.x.55
    roles:
...
...

```

OpenShift

- CEM-21614 In an OpenShift 4.6 cluster, the `contrail-status` command might show warnings for Zookeeper and RabbitMQ and status not reported. As a workaround, you can ignore this warning and based on other active contrail services, you can consider the Zookeeper and RabbitMQ statuses also as active. To get the router agent status, enter `contrail-status -t 15` on the compute node.
- CEM-20802 When creating any new user-defined namespace on Openshift-4.x/Contrail, by default SNAT is enabled and so all the pods part of this namespace by default can reach internet servers. As a workaround, explicitly configure the Contrail annotations on the namespace as `"opencontrail.org/ip_fabric_snat": "false"`.

DPDK and SR-IOV

- CEM-22835 During the high throughput data transfer between virtual machines (VMs) running in same DPDK enabled compute node, intermittent performance drop is observed along with checksum errors. Functionality is not impacted.
- CEM-21547 The SR-IOV configuration automation does not work for a Multihomed server.
- CEM-18922 On DPDK compute, memory of the VMs are mapped to only one numa. VM creation fails after the hugepages in that numa are exhausted if it is launched with `hw:mem_page_size='any'` flavor. As a workaround, use the `hw:mem_page_size='large'` flavor instead to avoid the issue.
- CEM-18408 In DPDK1911 with X710 NIC performance degrades due to mbuf leak if tx and rx are configured. Intel recommends configuring atleast 1K tx and rx descriptors on Fortville NICs for better and consistent performance, but they seem to have a degrading effect on X710 NIC.
- CEM-18163 On a DPDK compute, if `contrail-vrouter-agent` crashes or if `contrail-vrouter-agent` is restarted in a scaled setup with many sub-interfaces, all the sub-interfaces and their parent interface may become inactive. As a workaround, stop / start the instances whose interfaces are down.
- CEM-17883 VLAN tag does not work with Mellanox CX5 cards with DPDK 19.11.

- CEM-15561 vRouter offload with Mellanox NIC cards does not work. However the DPDK on Mellanox NICs without offload is supported.
- CEM-13685 DPDK vRouter with MLNX CX5 takes about 10 minutes and also lcore crash is seen. This happens once during initial installation.
- JCB-177787 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, jumbo MTU cannot be set. Maximum MTU allowed: <=1500.

SmartNIC

- CEM-20620 Gatewayless forwarding feature is not supported on Netronome.
- CEM-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.

General Routing

- CEM-21835 An allowed address pair (AAP) less than /32 must not be used, if Bare Metal Server (BMS) route (VXLAN route) is within the AAP CIDR.
- CEM-7262 You must not attach the Bidirectional Forwarding Detection (BFD) protocol a BGPaaS object. You can attach the BFD to to a VMI.
- CEM-21570 The AS_PATH retain does not work with a 4-Byte autonomous system number (ASN) enabled cluster.
- CEM-20421 In Contrail Networking, the logical router (LR) does not support dynamic next-hop port-mirroring, when Juniper headers are enabled. The Juniper header is not supported in port-mirroring as VXLAN is not the tunnel type used for the dynamic next hop in this case.
- CEM-20419 BFD session takes longer time to come up after the agent is restarted when MAC/IP is enabled in the VN and associated with BFD health check with target-ip set to "all".
- CEM-5043 VNI update on a LR doesn't update the RouteTable. As a workaround, delete the LogicalRouter and create a new LogicalRouter with the new VNI.
- CEM-3245 Multicast traffic originated from type-6 incapable QFX devices are duplicated by vRouters.
- CEM-22632 With an MX Series router acting as Datacenter Gateway and if MPLSoUDP is used between the controllers and MX, Floating IP use-case does not work. Use MPLSoGRE or VXLAN instead.

Contrail Fabric Management

- CEM-23931 Provisioning Contrail Fabric Manager using Ansible deployer intermittently fails while installing swift-client. If this condition occurs, preinstall swift-client on the controllers and then rerun the provisioning.
- CEM-20829 In Contrail fabric management with ML2 plugin, the telemetry_in_band_interface used for sFlow must be a physical interface. VLAN interfaces are not supported.
- CEM-18381 QFX5120 cannot be used as border leaf role in SP style for CRB role.
- CEM-15809 Updating VLAN-ID on a VPG in an enterprise style fabric is not supported. As a workaround, delete and recreate the fabric.
- CEM-14679 In fabric un-managed PNF use-case, some bogus static routes are pushed by DM under LR VRF on spines in case of CRB gateway.

As a workaround, change the value of the dummy_ip variable inside the device_manager docker. This line number below is based on the 2008 release code-base.

https://github.com/Juniper/contrail-controller/blob/R2008/src/config/fabric-ansible/ansible-playbooks/filter_plugins/fabric.py#L2594

After changing the value to the desired subnet and saving the file, restart the DM docker to reflect the change. Note that, this step should be performed at the beginning before fabric onboarding.

- CEM-14264 In release 2003, the Virtual Port Group create workflow will not pre-populate the VLAN-ID with the existing value that was defined with the first VPG for a given virtual network. The field is editable unlike in previous releases. This issue occurs in a fabric that was provisioned with the **Fabric-wide VLAN-ID significance** checkbox enabled.
- CEM-13767 Though Contrail fabric manager has the ability for the user to use custom image names for the fabric devices, for platforms like QFX10000-60C which runs on vmhost-based platforms, while uploading the image to CFM, the image name should be chosen in **junos-vmhost-install-x.tgz** format.
- CEM-8701 Onboarding of multiple BMS in parallel on SP-style fabric does not work. While bringing up a BMS using the Life Cycle Management workflow, sometimes on faster servers the re-image does not go through and instance not moved from ironic vn to tenant vn. This is because if the PXE boot request from the BMS is sent before the routes are converged between the BMS port and the TFTP service running in Contrail nodes. As a workaround, the servers can be rebooted or the BIOS in the servers can be configured to have a delayed boot.
- CEM-8149 BMS LCM with fabric set with enterprise_style=True is not supported. By default, enterprise_style is set to False. Avoid using enterprise_style=True if the fabric object onboards the BMS LCM instance.

- CEM-4358 In Contrail fabric deployments configuring QFX5110 as spine (CRB-Gateway) does not work.
- CEM-3959 BMS movement across TORs is not supported. To move BMS across TORs the whole VPG needs to be moved. That means if there are more than one BMS associated to one VPG, and one of the BMS need to be moved, the whole VPG need to be deleted and re-configured as per the new association.
- CEM-20794 In Contrail fabric management with ML2 plugin, configuring LAG in VMs associated with SRIOV VFs will not work.
- CEM-20693 The BGP routes widget under **Fabrics > Ports > *Leaf network device*** account for routes from inet.0 table only.
- CEM-20272 In L2 DCI mode, if selected fabrics have the same overlay ASN numbers, overlay iBGP is used between fabric devices for L2 DCI mode. In this case, the border device (physical router) marked with DCI-gateway RB role (routing and bridging) must also have RR (route reflector) RB role assigned. Without RR RB role, Overlay IBGP session won't stretch Layer 2 tenant virtual network across the fabric's leaf devices. So, we recommend that for L2 DCI Mode, ensure the physical router device is marked as DCI Gateway RB role along with RR role.
- CEM-19802 Security Groups cannot be used on QFX10K interfaces.

Kubernetes

- JCB-187287 High Availability provisioning of Kubernetes master is not supported.

General

- JCB-184776 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.

- CEM-4370 Additional links cannot be appended to service templates used to create PNF service chaining. If there is a need to add additional links, the service template needs to be deleted and re-added again.

- CEM-18398 Contrail WebUI doesn't work for System/Node status monitoring. As a workaround, check using CLI on the relevant nodes. This will not impact functionality.
- CEM-20414 During Contrail Command container restart one deploy job per contrail_cluster object is fired and leads to restart issues. As a workaround, modify the contrail-cluster or openstack-cluster object through commandcli instead of modifying/creating the endpoint directly through UI.
- CEM-17648 In case of BMS to BMS EVPN "Transparent" service chaining, Tunneled packet sent out of Transparent service instance to QFX have vlan-id and hence Traffic from left-bms to right-bms gets dropped since the inner header of the tunneled packet has vlan-id info which is internal to vRouter and QFX is not aware of the vlan-id so the packet gets dropped by the switch.
- CEM-17562 Under Security Groups, the entry appearing with **__no_rule__** can be ignored.
- CEM-16855 IPv6 ipam subnet option "enable_dhcp" is always ignored.

Telemetry and Analytics

- CEM-21526 When you upgrade Contrail Networking from a release prior to R2011 to Contrail Networking Release R2011.138, a Kafka container running on an analytics node might report connection issues.

Workaround: Perform the following steps:

1. Stop Kafka on all the analytics nodes.

```
docker stop analytics_alarm_kafka_1
```

2. On a single Contrail controller, perform the Zookeeper cleanup.

```
docker exec -it config_database_zookeeper_1 bash
```

```
bin/zkCli.sh -server <IP>:2181
```

```
deleteall /brokers
```

```
deleteall /consumers
```

3. Start Kafka on all the analytics nodes.

```
docker start analytics_alarm_kafka_1
```

- CEM-20846 In rare cases, sFlow node provisioning fails while initializing kafka container. If this scenario happens during provisioning, redeploying will bring up the sFlow nodes.
- CEM-20781 Telemetry KPI display for Junos EVO devices are not supported.

- CEM-18999 In a heavily scaled datacenter with around 128 racks and 4000 VNs, 256k VMIs, if the Contrail Insights OpeSstack adapter is restarted, it might take around 4 hrs for it to re-sync with the API server.
- CEM-13380 AppFormix Flows does not show up for multi homed devices on the fabric
- CEM-10929 When Contrail Insights is querying LLDP table from a device through SNMP, if SNMP calls time out, Contrail Insights marks the device as invalidConfiguration and notifies the user to take a look. When the user verifies that snmpwalk is working and there are no network issues, click **Edit** and reconfigure that device from **Settings > Network Devices** to make Contrail Insights try to run LLDP discovery and add this device again.

Known Behavior in Contrail Networking Release 2011.L2

- CEM-23142 Few deployment instances of RHOSP13 deployment with Contrail Networking 2011.L2 failed during installation of Openstack components. RHOSP13 customers are advised to contact Redhat support referring Case ID - 02995585 prior to installation.
- CEM-22783 Upgrading a RHOSP13 cluster from RHOSP13 images tested with Contrail Networking 2011.L1 to the RHOSP13 images tested with Contrail Networking 2011.L2 fails.

As a workaround, disable the Ceilometer service before upgrading. It is recommended to contact Redhat support and refer the case ID 02989025 before proceeding with upgrade.
- CEM-22835 During the high throughput data transfer between virtual machines (VMs) running in same DPDK enabled compute node, intermittent performance drop is observed along with checksum errors. Functionality is not impacted.
- CEM-22632 With an MX Series router acting as Datacenter Gateway and if MPLSoUDP is used between the controllers and MX, Floating IP use-case does not work. Use MPLSoGRE or VXLAN instead.
- CEM-21835 An allowed address pair (AAP) less than /32 must not be used, if Bare Metal Server (BMS) route (VXLAN route) is within the AAP CIDR.
- CEM-7262 You must not attach the Bidirectional Forwarding Detection (BFD) protocol a BGPaaS object. You can attach the BFD to to a VMI.
- CEM-21570 The AS_PATH retain does not work with a 4-Byte autonomous system number (ASN) enabled cluster.

- CEM-21614 In an OpenShift 4.6 cluster, the `contrail-status` command might show warnings for Zookeeper and RabbitMQ and status not reported. As a workaround, you can ignore this warning and based on other active contrail services, you can consider the Zookeeper and RabbitMQ statuses also as active. To get the router agent status, enter `contrail-status -t 15` on the compute node.
- CEM-21547 The SR-IOV configuration automation does not work for a Multihomed server.
- CEM-21526 When you upgrade Contrail Networking from a release prior to R2011 to Contrail Networking Release R2011.138, a Kafka container running on an analytics node might report connection issues.

Workaround: Perform the following steps:

1. Stop Kafka on all the analytics nodes.

```
docker stop analytics_alarm_kafka_1
```

2. On a single Contrail controller, perform the Zookeeper cleanup.

```
docker exec -it config_database_zookeeper_1 bash
```

```
bin/zkCli.sh -server <IP>:2181
```

```
deleteall /brokers
```

```
deleteall /consumers
```

3. Start Kafka on all the analytics nodes.

```
docker start analytics_alarm_kafka_1
```

- CEM-20846 In rare cases, sFlow node provisioning fails while initializing kafka container. If this scenario happens during provisioning, redeploying will bring up the sFlow nodes.
- CEM-20829 In Contrail fabric management with ML2 plugin, the `telemetry_in_band_interface` used for sFlow must be a physical interface. VLAN interfaces are not supported.
- CEM-20802 When creating any new user-defined namespace on Openshift-4.x/Contrail, by default SNAT is enabled and so all the pods part of this namespace by default can reach internet servers. As a workaround, explicitly configure the Contrail annotations on the namespace as `"opencontrail.org/ip_fabric_snat": "false"`.
- CEM-20794 In Contrail fabric management with ML2 plugin, configuring LAG in VMs associated with SRIOV VFs will not work.
- CEM-20781 Telemetry KPI display for Junos EVO devices are not supported.
- CEM-20693 The BGP routes widget under **Fabrics > Ports > *Leaf network device*** account for routes from inet.0 table only.

- CEM-20620 Gatewayless forwarding feature is not supported on Netronome.
- CEM-20421 In Contrail Networking, the logical router (LR) does not support dynamic next-hop port-mirroring, when Juniper headers are enabled. The Juniper header is not supported in port-mirroring as VXLAN is not the tunnel type used for the dynamic next hop in this case.
- CEM-20419 BFD session takes longer time to come up after the agent is restarted when MAC/IP is enabled in the VN and associated with BFD health check with target-ip set to "all".
- CEM-20414 During Contrail Command container restart one deploy job per contrail_cluster object is fired and leads to restart issues. As a workaround, modify the contrail-cluster or openstack-cluster object through commandcli instead of modifying/creating the endpoint directly through UI.
- CEM-20308 After one of HA Master nodes failover (or) vrouter restart, further user PODs creation might fail without getting IP-address.

As a workaround, find HA master nodes control-pods which are not in sync with respect to "new user-pod" and to restart them. Perform the following steps:

1. Log in to 3 HA masters and find the crictl pod with name "control" and log in to it verify the command output of `curl --cert /etc/certificates/server-key-localhost --insecure https://localhost:8083/Snh_IFMapTableShowReq?table_name=virtual-machine` showing the name of the latest user pod which failed.
 2. Restart those control PODs which are not in sync.
- CEM-20280 In a Kubernetes and OpenStack joint setup, vrouter-agent restart sometimes leads to an unauthorized operation error. To resolve the issue, restart the vrouter-agent again.
 - CEM-20272 In L2 DCI mode, if selected fabrics have the same overlay ASN numbers, overlay iBGP is used between fabric devices for L2 DCI mode. In this case, the border device (physical router) marked with DCI-gateway RB role (routing and bridging) must also have RR (route reflector) RB role assigned. Without RR RB role, Overlay iBGP session won't stretch Layer 2 tenant virtual network across the fabric's leaf devices. So, we recommend that for L2 DCI Mode, ensure the physical router device is marked as DCI Gateway RB role along with RR role.
 - CEM-19802 Security Groups cannot be used on QFX10K interfaces.
 - CEM-18979 The vRouter to vRouter encryption feature is beta quality and should be used for future product capability demonstrations only.
 - CEM-18999 In a heavily scaled datacenter with around 128 racks and 4000 VNs, 256k VMIs, if the Contrail Insights OpeSstack adapter is restarted, it might take around 4 hrs for it to re-sync with the API server.

- CEM-18922 On DPDK compute, memory of the VMs are mapped to only one numa. VM creation fails after the hugepages in that numa are exhausted if it is launched with `hw:mem_page_size='any'` flavor. As a workaround, use the `hw:mem_page_size='large'` flavor instead to avoid the issue.
- CEM-18408 In DPDK1911 with X710 NIC performance degrades due to mbuf leak if tx and rx are configured. Intel recommends configuring atleast 1K tx and rx descriptors on Fortville NICs for better and consistent performance, but they seem to have a degrading effect on X710 NIC.
- CEM-18398 Contrail WebUI doesn't work for System/Node status monitoring. As a workaround, check using CLI on the relevant nodes. This will not impact functionality.
- CEM-18381 QFX5120 cannot be used as border leaf role in SP style for CRB role.
- CEM-18163 On a DPDK compute, if contrail-vrouter-agent crashes or if contrail-vrouter-agent is restarted in a scaled setup with many sub-interfaces, all the sub-interfaces and their parent interface may become inactive. As a workaround, stop / start the instances whose interfaces are down.
- CEM-17991 In an OpenStack HA setup provisioned using Kolla and OpenStack Rocky, if you shut down all the servers at the same time and bring them up later, the Galera cluster fails. To recover the Galera cluster, follow these steps:
 1. Edit the `/etc/kolla/mariadb/galera.cnf` file to remove the `wsrep` address on one of the controllers as shown here.

```
wsrep_cluster_address = gcomm://
#wsrep_cluster_address = gcomm://10.x.x.8:4567,10.x.x.10:4567,10.x.x.11:4567
```

NOTE: If all the controllers are shut down in the managed scenario at the same time, you must select the controller that was shut down last.

2. Docker start mariadb on the controller on which you edited the file.
 3. Wait for a couple of minutes, ensure that the mariadb container is not restarting, and then Docker start mariadb on the remaining controllers.
 4. Restore the `/etc/kolla/mariadb/galera.cnf` file changes and restart the mariadb container on the previously selected controller.
- CEM-17883 VLAN tag does not work with Mellanox CX5 cards with DPDK 19.11.
 - CEM-17648 In case of BMS to BMS EVPN “Transparent” service chaining, Tunneled packet sent out of Transparent service instance to QFX have vlan-id and hence Traffic from left-bms to right-bms gets

dropped since the inner header of the tunneled packet has vlan-id info which is internal to vRouter and QFX is not aware of the vlan-id so the packet gets dropped by the switch.

- CEM-17562 Under Security Groups, the entry appearing with **__no_rule__** can be ignored.
- CEM-16855 IPv6 ipam subnet option "enable_dhcp" is always ignored.
- CEM-15809 Updating VLAN-ID on a VPG in an enterprise style fabric is not supported. As a workaround, delete and recreate the fabric.
- CEM-15764 In Octavia Load Balancer, traffic destined to the Floating IP of the load balancer VM does not get directed to the backend VMs. Traffic destined to the actual VM IP of the Load Balancer VM will work fine.
- CEM-15561 vRouter offload with Mellanox NIC cards does not work. However the DPDK on Mellanox NICs without offload is supported.
- CEM-14679 In fabric un-managed PNF use-case, some bogus static routes are pushed by DM under LR VRF on spines in case of CRB gateway.

As a workaround, change the value of the dummy_ip variable inside the device_manager docker. This line number below is based on the 2008 release code-base.

https://github.com/Juniper/contrail-controller/blob/R2008/src/config/fabric-ansible/ansible-playbooks/filter_plugins/fabric.py#L2594

After changing the value to the desired subnet and saving the file, restart the DM docker to reflect the change. Note that, this step should be performed at the beginning before fabric onboarding.

- CEM-14264 In release 2003, the Virtual Port Group create workflow will not pre-populate the VLAN-ID with the existing value that was defined with the first VPG for a given virtual network. The field is editable unlike in previous releases. This issue occurs in a fabric that was provisioned with the **Fabric-wide VLAN-ID significance** checkbox enabled.
- CEM-13767 Though Contrail fabric manager has the ability for the user to use custom image names for the fabric devices, for platforms like QFX10000-60C which runs on vmhost-based platforms, while uploading the image to CFM, the image name should be chosen in **junos-vmhost-install-x.tgz** format.
- CEM-13685 DPDK vRouter with MLNX CX5 takes about 10 minutes and also lcore crash is seen. This happens once during initial installation.
- CEM-13380 AppFormix Flows does not show up for multi homed devices on the fabric
- CEM-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.

- CEM-10929 When Contrail Insights is querying LLDP table from a device through SNMP, if SNMP calls time out, Contrail Insights marks the device as invalidConfiguration and notifies the user to take a look. When the user verifies that snmpwalk is working and there are no network issues, click **Edit** and reconfigure that device from **Settings > Network Devices** to make Contrail Insights try to run LLDP discovery and add this device again.
- CEM-9979 During upgrade of DPDK computes deployed with OOO Heat Templates in RHOSP environment, vRouter coredumps are observed. This is due to the sequence in which the services are started during upgrade and does not have impact on cluster operation.
- CEM-8701 Onboarding of multiple BMS in parallel on SP-style fabric does not work. While bringing up a BMS using the Life Cycle Management workflow, sometimes on faster servers the re-image does not go through and instance not moved from ironic vn to tenant vn. This is because if the PXE boot request from the BMS is sent before the routes are converged between the BMS port and the TFTP service running in Contrail nodes. As a workaround, the servers can be rebooted or the BIOS in the servers can be configured to have a delayed boot.
- CEM-8149 BMS LCM with fabric set with enterprise_style=True is not supported. By default, enterprise_style is set to False. Avoid using enterprise_style=True if the fabric object onboards the BMS LCM instance.
- CEM-5141 For deleting compute nodes, the UI workflow will not work. Instead, update the instances.yaml with "ENABLE_DESTROY: True" and "roles:" (leave it empty) and run the following playbooks.

```
ansible-playbook -i inventory/ -e orchestrator=openstack --tags nova playbooks/
install_openstack.yml
ansible-playbook -i inventory/ -e orchestrator=openstack playbooks/install_contrail.yml
```

For example:

```
global_configuration:
  ENABLE_DESTROY: True
...
...
instances:
...
...
  srvr5:
    provider: bms
    ip: 19x.xxx.x.55
    roles:
```

...

...

- CEM-5043 VNI update on a LR doesn't update the RouteTable. As a workaround, delete the LogicalRouter and create a new LogicalRouter with the new VNI.
- CEM-4370 Additional links cannot be appended to service templates used to create PNF service chaining. If there is a need to add additional links, the service template needs to be deleted and re-added again.
- CEM-4358 In Contrail fabric deployments configuring QFX5110 as spine (CRB-Gateway) does not work.
- CEM-3959 BMS movement across TORs is not supported. To move BMS across TORs the whole VPG needs to be moved. That means if there are more than one BMS associated to one VPG, and one of the BMS need to be moved, the whole VPG need to be deleted and re-configured as per the new association.
- CEM-3245 Multicast traffic originated from type-6 incapable QFX devices are duplicated by vRouters.
- JCB-187287 High Availability provisioning of Kubernetes master is not supported.
- JCB-184776 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.

- JCB-177787 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, jumbo MTU cannot be set. Maximum MTU allowed: ≤ 1500 .

Known Behavior in Contrail Networking Release 2011.L1

- CEM-21910 Post deployment of a RHOSP16 cluster with IPA enabled, sometimes the vRouter agent will be in the initializing state. To clear this and bring back the agent to normal, issue the following command on the compute node:

```
sudo podman restart
contrail_vrouter_agent
```

- CEM-21909 The deployment of a RHOSP16 cluster with IPA enabled needs the following parameters to be set in any heat-templates used for deployment:

```
parameter_defaults:
  InternalTLSVncCAFile: /etc/ipa/ca.crt
  InternalTLSQemuCAFile: /etc/ipa/ca.crt
  InternalTLSVncProxyCAFile: /etc/ipa/ca.crt
```

- CEM-21835 An allowed address pair (AAP) less than /32 must not be used, if Bare Metal Server (BMS) route (VXLAN route) is within the AAP CIDR.
- CEM-21793 A BGP multihop session cannot be established with a Kubernetes container via a trunked subport when the Kubernetes container is running inside a VM instance that is managed by Openstack because Contrail Networking responds to the BGP request with the wrong MAC address.
- CEM-21446 If the sFlow node deletion fails for the first time, re-provision the cluster to remove the node. The other nodes will stay active during the reprovisioning.
- CEM-7262 You must not attach the Bidirectional Forwarding Detection (BFD) protocol a BGPaaS object. You can attach the BFD to to a VMI.
- CEM-21570 The AS_PATH retain does not work with a 4-Byte autonomous system number (ASN) enabled cluster.
- CEM-21614 In an OpenShift 4.6 cluster, the contrail-status command might show warnings for Zookeeper and RabbitMQ and status not reported. As a workaround, you can ignore this warning and based on other active contrail services, you can consider the Zookeeper and RabbitMQ statuses also as active. To get the router agent status, enter `contrail-status -t 15` on the compute node.
- CEM-21547 The SR-IOV configuration automation does not work for a Multihomed server.

- CEM-21526 When you upgrade Contrail Networking from a release prior to R2011 to Contrail Networking Release R2011.138, a Kafka container running on an analytics node might report connection issues.

Workaround: Perform the following steps:

1. Stop Kafka on all the analytics nodes.

```
docker stop analytics_alarm_kafka_1
```

2. On a single Contrail controller, perform the Zookeeper cleanup.

```
docker exec -it config_database_zookeeper_1 bash
```

```
bin/zkCli.sh -server <IP>:2181
```

```
deleteall /brokers
```

```
deleteall /consumers
```

3. Start Kafka on all the analytics nodes.

```
docker start analytics_alarm_kafka_1
```

- CEM-20846 In rare cases, sFlow node provisioning fails while initializing kafka container. If this scenario happens during provisioning, redeploying will bring up the sFlow nodes.
- CEM-20829 In Contrail fabric management with ML2 plugin, the `telemetry_in_band_interface` used for sFlow must be a physical interface. VLAN interfaces are not supported.
- CEM-20828 In Contrail fabric management with ML2 plugin, the sFlow nodes are referred to with the UUID and cannot be referred to by the hostname.
- CEM-20802 When creating any new user-defined namespace on Openshift-4.x/Contrail, by default SNAT is enabled and so all the pods part of this namespace by default can reach internet servers. As a workaround, explicitly configure the Contrail annotations on the namespace as `"opencontrail.org/ip_fabric_snat": "false"`.
- CEM-20794 In Contrail fabric management with ML2 plugin, configuring LAG in VMs associated with SRIOV VFs will not work.
- CEM-20781 Telemetry KPI display for Junos EVO devices are not supported.
- CEM-20693 The BGP routes widget under **Fabrics > Ports > Leaf network device** account for routes from inet.0 table only.
- CEM-20620 Gatewayless forwarding feature is not supported on Netronome.
- CEM-20513 While upgrading Contrail Networking Release 19xx with RHOSP13 to Contrail Networking Release 2011 with RHOSP16.1, FFU upgrade of compute node fails with the error—'RPC

failed at server. Insufficient access: Insufficient 'add' privilege to add the entry 'krbprincipalname=qemu/compute-0-ffu.internalapi.nodel8.local@NODEL8.LOCAL,cn=services,cn=accounts,dc=nodel8,dc=local'.

Perform the following steps to apply the patch:

- On undercloud:
 1. Apply the patch as given at <https://review.opendev.org/c/openstack/tripleo-heat-templates/+/-/764064/3/deployment/nova/novajoin-container-puppet.yaml>.

2. Deploy undercloud by running `openstack undercloud install` command.

3. Ensure the novajoin_notifier container is up and not constantly restarting.

```
podman ps |grep novajoin
```

Validate the status by checking `/var/log/containers/novajoin/*` logs.

TLS-E can not work properly without the novajoin containers running.

4. Ensure that the `transport_url` for novajoin does not contain guest user.

```
grep ^transport_url /var/lib/config-data/puppet-generated/novajoin/etc/novajoin/join.conf
```

- On compute (kernel/dpdk) node:

1. Log in to knit admin.

```
heat-admin@overcloud-contraildpdk-0-ffu ~]$ kinit admin
```

Use FreeIPA password for `admin@NODEL8.LOCAL`.

2. Add DNS recorded of compute node.

```
[heat-admin@overcloud-contraildpdk-0-ffu ~]$ ipa dnsrecord-add
Record name: overcloud-contraildpdk-0-ffu
Zone name: internalapi.nodel8.local
Please choose a type of DNS resource record to be added
The most common types for this type of zone are: A, AAAA
DNS resource record type: a
A IP Address: 192.168.xx.18
Record name: overcloud-contraildpdk-0-ffu
A record: 192.168.xx.18 #( IP address of compute node from openstack server list)
```

3. Add IPA service.

```
[heat-admin@overcloud-contraildpdk-0-ffu ~]$ ipa service-add qemu/overcloud-
contraildpdk-0-ffu.internalapi.node18.local@NODE18.LOCAL
Principal name: qemu/overcloud-contraildpdk-0-ffu.internalapi.node18.local@NODE18.LOCAL
Principal alias: qemu/overcloud-contraildpdk-0-
ffu.internalapi.node18.local@NODE18.LOCAL
Managed by: overcloud-contraildpdk-0-ffu.internalapi.node18.local
```

4. Add host to the IPA service.

```
[heat-admin@overcloud-contraildpdk-0-ffu ~]$ ipa service-add-host --hosts overcloud-
contraildpdk-0-ffu.node18.local qemu/overcloud-contraildpdk-0-
ffu.internalapi.node18.local@NODE18.LOCAL
Principal name: qemu/overcloud-contraildpdk-0-ffu.internalapi.node18.local@NODE18.LOCAL
Principal alias: qemu/overcloud-contraildpdk-0-
ffu.internalapi.node18.local@NODE18.LOCAL
Managed by: overcloud-contraildpdk-0-ffu.internalapi.node18.local, overcloud-
contraildpdk-0-ffu.node18.local
Number of members added 1
```

- CEM-20421 In Contrail Networking, the logical router (LR) does not support dynamic next-hop port-mirroring, when Juniper headers are enabled. The Juniper header is not supported in port-mirroring as VXLAN is not the tunnel type used for the dynamic next hop in this case.
- CEM-20419 BFD session takes longer time to come up after the agent is restarted when MAC/IP is enabled in the VN and associated with BFD health check with target-ip set to "all".
- CEM-20414 During Contrail Command container restart one deploy job per contrail_cluster object is fired and leads to restart issues. As a workaround, modify the contrail-cluster or openstack-cluster object through commandcli instead of modifying/creating the endpoint directly through UI.
- CEM-20308 After one of HA Master nodes failover (or) router restart, further user PODs creation might fail without getting IP-address.

As a workaround, find HA master nodes control-pods which are not in sync with respect to "new user-pod" and to restart them. Perform the following steps:

1. Log in to 3 HA masters and find the crictl pod with name "control" and log in to it verify the command output of "curl --cert /etc/certificates/server-key-localhost --insecure https://localhost:8083/Snh_IFMapTableShowReq?table_name=virtual-machine" showing the name of the latest user pod which failed.

2. Restart those control PODs which are not in sync.

- CEM-20280 In a Kubernetes and OpenStack joint setup, vrouter-agent restart sometimes leads to an unauthorized operation error. To resolve the issue, restart the vrouter-agent again.
- CEM-20272 In L2 DCI mode, if selected fabrics have the same overlay ASN numbers, overlay iBGP is used between fabric devices for L2 DCI mode. In this case, the border device (physical router) marked with DCI-gateway RB role (routing and bridging) must also have RR (route reflector) RB role assigned. Without RR RB role, Overlay iBGP session won't stretch Layer 2 tenant virtual network across the fabric's leaf devices. So, we recommend that for L2 DCI Mode, ensure the physical router device is marked as DCI Gateway RB role along with RR role.
- CEM-20248 While upgrading Contrail Networking Release 19xx with RHOSP13 to Contrail Networking Release 2011 with RHOSP16.1, overcloud nodes transitions to ERROR state after upgrading the undercloud. As a workaround, apply a patch as mentioned at https://bugzilla.redhat.com/show_bug.cgi?id=1850929.
- CEM-19802 Security Groups cannot be used on QFX10K interfaces.
- CEM-19151 During deployment we see race condition, due to which ipa-client installation on compute nodes fails. This is an issue with Red Hat. As a workaround, before deployment starts, modify the following file to add sleep of 400 seconds on undercloud.

```

name: DBG
    sudo vi /usr/share/ansible/roles/tripleo-kernel/tasks/kernelargs.yml -
    debug:
        msg: "sleep 400 sec if reboot_required == {{ reboot_required }}"
    - name: DBG sleep
      shell: sleep 400
    when:
        - reboot_required is defined and reboot_required# then
    sudo find / -name kernelargs.yml
    # to find all such files on undercloud and in containers because I am not sure
    which exactly is used (from host or from container)
    # and overwriting such files in containers like
    sudo cp /usr/share/ansible/roles/tripleo-kernel/tasks/
kernelargs.yml      /var/lib/containers/storage/overlay/
6dc6b96b1392e5302b63156fa093525e17131bef1203cad005a911ad09241f5a/diff/usr/share/ansible/roles/
tripleo-kernel/tasks/kernelargs.yml

```

- CEM-18979 The vRouter to vRouter encryption feature is beta quality and should be used for future product capability demonstrations only.

- CEM-18999 In a heavily scaled datacenter with around 128 racks and 4000 VNs, 256k VMIs, if the Contrail Insights OpeSstack adapter is restarted, it might take around 4 hrs for it to re-sync with the API server.
- CEM-18922 On DPDK compute, memory of the VMs are mapped to only one numa. VM creation fails after the hugepages in that numa are exhausted if it is launched with `hw:mem_page_size='any'` flavor. As a workaround, use the `hw:mem_page_size='large'` flavor instead to avoid the issue.
- CEM-18909 In case of RHOSP16 deployment with TLS, XMPP connection down is seen post deployment completion. While this is a cosmetic issue and does not impact functionality, as a workaround, restart the vRouter agent container on all compute nodes to update status.
- CEM-18408 In DPDK1911 with X710 NIC performance degrades due to mbuf leak if tx and rx are configured. Intel recommends configuring atleast 1K tx and rx descriptors on Fortville NICs for better and consistent performance, but they seem to have a degrading effect on X710 NIC.
- CEM-18398 Contrail WebUI doesn't work for System/Node status monitoring. As a workaround, check using CLI on the relevant nodes. This will not impact functionality.
- CEM-18381 QFX5120 cannot be used as border leaf role in SP style for CRB role.
- CEM-18163 On a DPDK compute, if contrail-vrouter-agent crashes or if contrail-vrouter-agent is restarted in a scaled setup with many sub-interfaces, all the sub-interfaces and their parent interface may become inactive. As a workaround, stop / start the instances whose interfaces are down.
- CEM-17991 In an OpenStack HA setup provisioned using Kolla and OpenStack Rocky, if you shut down all the servers at the same time and bring them up later, the Galera cluster fails. To recover the Galera cluster, follow these steps:
 1. Edit the `/etc/kolla/mariadb/galera.cnf` file to remove the `wsrep` address on one of the controllers as shown here.

```
wsrep_cluster_address = gcomm://
#wsrep_cluster_address = gcomm://10.x.x.8:4567,10.x.x.10:4567,10.x.x.11:4567
```

NOTE: If all the controllers are shut down in the managed scenario at the same time, you must select the controller that was shut down last.

2. Docker start mariadb on the controller on which you edited the file.
3. Wait for a couple of minutes, ensure that the mariadb container is not restarting, and then Docker start mariadb on the remaining controllers.

4. Restore the `/etc/kolla/mariadb/galera.cnf` file changes and restart the mariadb container on the previously selected controller.

- CEM-17883 VLAN tag does not work with Mellanox CX5 cards with DPDK 19.11.
- CEM-17866 Monitoring/Operations page crashes with "Cannot read property 'className' of undefined". As a workaround, refresh the page to display the content properly.
- CEM-17648 In case of BMS to BMS EVPN "Transparent" service chaining, Tunneled packet sent out of Transparent service instance to QFX have vlan-id and hence Traffic from left-bms to right-bms gets dropped since the inner header of the tunneled packet has vlan-id info which is internal to vRouter and QFX is not aware of the vlan-id so the packet gets dropped by the switch.
- CEM-17562 Under Security Groups, the entry appearing with `__no_rule__` can be ignored.
- CEM-16855 IPv6 ipam subnet option "enable_dhcp" is always ignored.
- CEM-15809 Updating VLAN-ID on a VPG in an enterprise style fabric is not supported. As a workaround, delete and recreate the fabric.
- CEM-15764 In Octavia Load Balancer, traffic destined to the Floating IP of the load balancer VM does not get directed to the backend VMs. Traffic destined to the actual VM IP of the Load Balancer VM will work fine.
- CEM-15561 vRouter offload with Mellanox NIC cards does not work. However the DPDK on Mellanox NICs without offload is supported.
- CEM-14679 In fabric un-managed PNF use-case, some bogus static routes are pushed by DM under LR VRF on spines in case of CRB gateway.

As a workaround, change the value of the `dummy_ip` variable inside the `device_manager` docker. This line number below is based on the 2008 release code-base.

https://github.com/Juniper/contrail-controller/blob/R2008/src/config/fabric-ansible/ansible-playbooks/filter_plugins/fabric.py#L2594

After changing the value to the desired subnet and saving the file, restart the DM docker to reflect the change. Note that, this step should be performed at the beginning before fabric onboarding.

- CEM-14264 In release 2003, the Virtual Port Group create workflow will not pre-populate the VLAN-ID with the existing value that was defined with the first VPG for a given virtual network. The field is editable unlike in previous releases. This issue occurs in a fabric that was provisioned with the **Fabric-wide VLAN-ID significance** checkbox enabled.
- CEM-13767 Though Contrail fabric manager has the ability for the user to use custom image names for the fabric devices, for platforms like QFX10000-60C which runs on vmhost-based platforms, while uploading the image to CFM, the image name should be chosen in **junos-vmhost-install-x.tgz** format.

- CEM-13685 DPDK vRouter with MLNX CX5 takes about 10 minutes and also lcore crash is seen. This happens once during initial installation.
- CEM-13380 AppFormix Flows does not show up for multi homed devices on the fabric
- CEM-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.
- CEM-10929 When Contrail Insights is querying LLDP table from a device through SNMP, if SNMP calls time out, Contrail Insights marks the device as invalidConfiguration and notifies the user to take a look. When the user verifies that snmpwalk is working and there are no network issues, click **Edit** and reconfigure that device from **Settings > Network Devices** to make Contrail Insights try to run LLDP discovery and add this device again.
- CEM-9979 During upgrade of DPDK computes deployed with OOO Heat Templates in RHOSP environment, vRouter coredumps are observed. This is due to the sequence in which the services are started during upgrade and does not have impact on cluster operation.
- CEM-8701 Onboarding of multiple BMS in parallel on SP-style fabric does not work. While bringing up a BMS using the Life Cycle Management workflow, sometimes on faster servers the re-image does not go through and instance not moved from ironic vn to tenant vn. This is because if the PXE boot request from the BMS is sent before the routes are converged between the BMS port and the TFTP service running in Contrail nodes. As a workaround, the servers can be rebooted or the BIOS in the servers can be configured to have a delayed boot.
- CEM-8149 BMS LCM with fabric set with enterprise_style=True is not supported. By default, enterprise_style is set to False. Avoid using enterprise_style=True if the fabric object onboards the BMS LCM instance.
- CEM-5141 For deleting compute nodes, the UI workflow will not work. Instead, update the instances.yaml with "ENABLE_DESTROY: True" and "roles:" (leave it empty) and run the following playbooks.

```
ansible-playbook -i inventory/ -e orchestrator=openstack --tags nova playbooks/
install_openstack.yml
ansible-playbook -i inventory/ -e orchestrator=openstack playbooks/install_contrail.yml
```

For example:

```
global_configuration:
  ENABLE_DESTROY: True
  ...
  ...
instances:
```

```

...
...
  srvr5:
    provider: bms
    ip: 19x.xxx.x.55
    roles:
...
...

```

- CEM-5043 VNI update on a LR doesn't update the RouteTable. As a workaround, delete the LogicalRouter and create a new LogicalRouter with the new VNI.
- CEM-4370 Additional links cannot be appended to service templates used to create PNF service chaining. If there is a need to add additional links, the service template needs to be deleted and re-added again.
- CEM-4358 In Contrail fabric deployments configuring QFX5110 as spine (CRB-Gateway) does not work.
- CEM-3959 BMS movement across TORs is not supported. To move BMS across TORs the whole VPG needs to be moved. That means if there are more than one BMS associated to one VPG, and one of the BMS need to be moved, the whole VPG need to be deleted and re-configured as per the new association.
- CEM-3245 Multicast traffic originated from type-6 incapable QFX devices are duplicated by vRouters.
- JCB-187287 High Availability provisioning of Kubernetes master is not supported.
- JCB-184776 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.

- JCB-177787 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, jumbo MTU cannot be set. Maximum MTU allowed: <=1500.

Known Behavior in Contrail Networking Release 2011

- CEM-20846 In rare cases, sFlow node provisioning fails while initializing kafka container. If this scenario happens during provisioning, redeploying will bring up the sFlow nodes.
- CEM-20829 In Contrail fabric management with ML2 plugin, the `telemetry_in_band_interface` used for sFlow must be a physical interface. VLAN interfaces are not supported.
- CEM-20828 In Contrail fabric management with ML2 plugin, the sFlow nodes are referred to with the UUID and cannot be referred to by the hostname.
- CEM-20802 When creating any new user-defined namespace on Openshift-4.x/Contrail, by default SNAT is enabled and so all the pods part of this namespace by default can reach internet servers. As a workaround, explicitly configure the Contrail annotations on the namespace as `"opencontrail.org/ip_fabric_snat": "false"`.
- CEM-20794 In Contrail fabric management with ML2 plugin, configuring LAG in VMs associated with SRIOV VFs will not work.
- CEM-20781 Telemetry KPI display for Junos EVO devices are not supported.
- CEM-20693 The BGP routes widget under **Fabrics > Ports > *Leaf network device*** account for routes from inet.0 table only.
- CEM-20620 Gatewayless forwarding feature is not supported on Netronome.
- CEM-20513 While upgrading Contrail Networking Release 19xx with RHOSP13 to Contrail Networking Release 2011 with RHOSP16.1, FFU upgrade of compute node fails with the error—'RPC failed at server. Insufficient access: Insufficient 'add' privilege to add the entry 'krbprincipalname=qemu/compute-0-ffu.internalapi.node18.local@NODE18.LOCAL,cn=services,cn=accounts,dc=node18,dc=local'.

Perform the following steps to apply the patch:

- On undercloud:
 1. Apply the patch as given at <https://review.opendev.org/c/openstack/tripleo-heat-templates/+/-/764064/3/deployment/nova/novajoin-container-puppet.yaml>.
 2. Deploy undercloud by running `openstack undercloud install` command.
 3. Ensure the `novajoin_notifier` container is up and not constantly restarting.

```
podman ps |grep novajoin
```

Validate the status by checking `/var/log/containers/novajoin/*` logs.

TLS-E can not work properly without the novajoin containers running.

4. Ensure that the `transport_url` for novajoin does not contain guest user.

```
grep ^transport_url /var/lib/config-data/puppet-generated/novajoin/etc/novajoin/join.conf
```

- On compute (kernel/dpdk) node:

1. Log in to knit admin.

```
heat-admin@overcloud-contraildpdk-0-ffu ~]$ kinit admin
```

Use FreeIPA password for `admin@NODEL8.LOCAL`.

2. Add DNS recorded of compute node.

```
[heat-admin@overcloud-contraildpdk-0-ffu ~]$ ipa dnsrecord-add
Record name: overcloud-contraildpdk-0-ffu
Zone name: internalapi.nodel8.local
Please choose a type of DNS resource record to be added
The most common types for this type of zone are: A, AAAA
DNS resource record type: a
A IP Address: 192.168.xx.18
Record name: overcloud-contraildpdk-0-ffu
A record: 192.168.xx.18 #( IP address of compute node from openstack server list)
```

3. Add IPA service.

```
[heat-admin@overcloud-contraildpdk-0-ffu ~]$ ipa service-add qemu/overcloud-
contraildpdk-0-ffu.internalapi.nodel8.local@NODEL8.LOCAL
Principal name: qemu/overcloud-contraildpdk-0-ffu.internalapi.nodel8.local@NODEL8.LOCAL
Principal alias: qemu/overcloud-contraildpdk-0-
ffu.internalapi.nodel8.local@NODEL8.LOCAL
Managed by: overcloud-contraildpdk-0-ffu.internalapi.nodel8.local
```

4. Add host to the IPA service.

```
[heat-admin@overcloud-contraildpdk-0-ffu ~]$ ipa service-add-host --hosts overcloud-
contraildpdk-0-ffu.nodel8.local qemu/overcloud-contraildpdk-0-
ffu.internalapi.nodel8.local@NODEL8.LOCAL
Principal name: qemu/overcloud-contraildpdk-0-ffu.internalapi.nodel8.local@NODEL8.LOCAL
```

```
Principal alias: qemu/overcloud-contraildpdk-0-
ffu.internalapi.nodel8.local@NODEL8.LOCAL
Managed by: overcloud-contraildpdk-0-ffu.internalapi.nodel8.local, overcloud-
contraildpdk-0-ffu.nodel8.local
Number of members added 1
```

- CEM-20421 In Contrail Networking, the logical router (LR) does not support dynamic next-hop port-mirroring, when Juniper headers are enabled. The Juniper header is not supported in port-mirroring as VXLAN is not the tunnel type used for the dynamic next hop in this case.
- CEM-20419 BFD session takes longer time to come up after the agent is restarted when MAC/IP is enabled in the VN and associated with BFD health check with target-ip set to "all".
- CEM-20414 During Contrail Command container restart one deploy job per contrail_cluster object is fired and leads to restart issues. As a workaround, modify the contrail-cluster or openstack-cluster object through commandcli instead of modifying/creating the endpoint directly through UI.
- CEM-20308 After one of HA Master nodes failover (or) vrouter restart, further user PODs creation might fail without getting IP-address.

As a workaround, find HA master nodes control-pods which are not in sync with respect to "new user-pod" and to restart them. Perform the following steps:

1. Log in to 3 HA masters and find the crictl pod with name "control" and log in to it verify the command output of "curl --cert /etc/certificates/server-key-localhost --insecure https://localhost:8083/Snh_IFMapTableShowReq?table_name=virtual-machine" showing the name of the latest user pod which failed.

2. Restart those control PODs which are not in sync.

- CEM-20280 In a Kubernetes and OpenStack joint setup, vrouter-agent restart sometimes leads to an unauthorized operation error. To resolve the issue, restart the vrouter-agent again.
- CEM-20272 In L2 DCI mode, if selected fabrics have the same overlay ASN numbers, overlay iBGP is used between fabric devices for L2 DCI mode. In this case, the border device (physical router) marked with DCI-gateway RB role (routing and bridging) must also have RR (route reflector) RB role assigned. Without RR RB role, Overlay iBGP session won't stretch Layer 2 tenant virtual network across the fabric's leaf devices. So, we recommend that for L2 DCI Mode, ensure the physical router device is marked as DCI Gateway RB role along with RR role.
- CEM-20248 While upgrading Contrail Networking Release 19xx with RHOSP13 to Contrail Networking Release 2011 with RHOSP16.1, overcloud nodes transition to ERROR state after upgrading the undercloud. As a workaround, apply a patch as mentioned at https://bugzilla.redhat.com/show_bug.cgi?id=1850929.
- CEM-19802 Security Groups cannot be used on QFX10K interfaces.

- CEM-19151 During deployment we see race condition, due to which ipa-client installation on compute nodes fails. This is an issue with Red Hat. As a workaround, before deployment starts, modify the following file to add sleep of 400 seconds on undercloud.

```

sudo vi /usr/share/ansible/roles/tripleo-kernel/tasks/kernelargs.yml -
name: DBG
debug:
  msg: "sleep 400 sec if reboot_required == {{ reboot_required }}"
  - name: DBG sleep
    shell: sleep 400
  when:
    - reboot_required is defined and reboot_required# then
sudo find / -name kernelargs.yml
# to find all such files on undercloud and in containers because I am not sure
which exactly is used (from host or from container)
# and overwriting such files in containers like
sudo cp /usr/share/ansible/roles/tripleo-kernel/tasks/
kernelargs.yml /var/lib/containers/storage/overlay/
6dc6b96b1392e5302b63156fa093525e17131bef1203cad005a911ad09241f5a/diff/usr/share/ansible/roles/
tripleo-kernel/tasks/kernelargs.yml

```

- CEM-18979 The vRouter to vRouter encryption feature is beta quality and should be used for future product capability demonstrations only.
- CEM-18999 In a heavily scaled datacenter with around 128 racks and 4000 VNs, 256k VMIs, if the Contrail Insights OpeSstack adapter is restarted, it might take around 4 hrs for it to re-sync with the API server.
- CEM-18922 On DPDK compute, memory of the VMs are mapped to only one numa. VM creation fails after the hugepages in that numa are exhausted if it is launched with `hw:mem_page_size='any'` flavor. As a workaround, use the `hw:mem_page_size='large'` flavor instead to avoid the issue.
- CEM-18909 In case of RHOSP16 deployment with TLS, XMPP connection down is seen post deployment completion. While this is a cosmetic issue and does not impact functionality, as a workaround, restart the vRouter agent container on all compute nodes to update status.
- CEM-18408 In DPDK1911 with X710 NIC performance degrades due to mbuf leak if txd and rxd are configured. Intel recommends configuring atleast 1K tx and rx descriptors on Fortville NICs for better and consistent performance, but they seem to have a degrading effect on X710 NIC.
- CEM-18398 Contrail WebUI doesn't work for System/Node status monitoring. As a workaround, check using CLI on the relevant nodes. This will not impact functionality.
- CEM-18381 QFX5120 cannot be used as border leaf role in SP style for CRB role.

- CEM-18163 On a DPDK compute, if contrail-vrouter-agent crashes or if contrail-vrouter-agent is restarted in a scaled setup with many sub-interfaces, all the sub-interfaces and their parent interface may become inactive. As a workaround, stop / start the instances whose interfaces are down.
- CEM-17991 In an OpenStack HA setup provisioned using Kolla and OpenStack Rocky, if you shut down all the servers at the same time and bring them up later, the Galera cluster fails. To recover the Galera cluster, follow these steps:
 1. Edit the `/etc/kolla/mariadb/galera.cnf` file to remove the `wsrep` address on one of the controllers as shown here.

```
wsrep_cluster_address = gcomm://
#wsrep_cluster_address = gcomm://10.x.x.8:4567,10.x.x.10:4567,10.x.x.11:4567
```

NOTE: If all the controllers are shut down in the managed scenario at the same time, you must select the controller that was shut down last.

2. Docker start mariadb on the controller on which you edited the file.
 3. Wait for a couple of minutes, ensure that the mariadb container is not restarting, and then Docker start mariadb on the remaining controllers.
 4. Restore the `/etc/kolla/mariadb/galera.cnf` file changes and restart the mariadb container on the previously selected controller.
- CEM-17883 VLAN tag does not work with Mellanox CX5 cards with DPDK 19.11.
 - CEM-17866 Monitoring/Operations page crashes with "Cannot read property 'className' of undefined". As a workaround, refresh the page to display the content properly.
 - CEM-17648 In case of BMS to BMS EVPN "Transparent" service chaining, Tunneled packet sent out of Transparent service instance to QFX have vlan-id and hence Traffic from left-bms to right-bms gets dropped since the inner header of the tunneled packet has vlan-id info which is internal to vRouter and QFX is not aware of the vlan-id so the packet gets dropped by the switch.
 - CEM-17562 Under Security Groups, the entry appearing with `__no_rule__` can be ignored.
 - CEM-16855 IPv6 ipam subnet option "enable_dhcp" is always ignored.
 - CEM-15809 Updating VLAN-ID on a VPG in an enterprise style fabric is not supported. As a workaround, delete and recreate the fabric.

- CEM-15764 In Octavia Load Balancer, traffic destined to the Floating IP of the load balancer VM does not get directed to the backend VMs. Traffic destined to the actual VM IP of the Load Balancer VM will work fine.
- CEM-15561 vRouter offload with Mellanox NIC cards does not work. However the DPDK on Mellanox NICs without offload is supported.
- CEM-14679 In fabric un-managed PNF use-case, some bogus static routes are pushed by DM under LR VRF on spines in case of CRB gateway.

As a workaround, change the value of the dummy_ip variable inside the device_manager docker. This line number below is based on the 2008 release code-base.

https://github.com/Juniper/contrail-controller/blob/R2008/src/config/fabric-ansible/ansible-playbooks/filter_plugins/fabric.py#L2594

After changing the value to the desired subnet and saving the file, restart the DM docker to reflect the change. Note that, this step should be performed at the beginning before fabric onboarding.

- CEM-14264 In release 2003, the Virtual Port Group create workflow will not pre-populate the VLAN-ID with the existing value that was defined with the first VPG for a given virtual network. The field is editable unlike in previous releases. This issue occurs in a fabric that was provisioned with the **Fabric-wide VLAN-ID significance** checkbox enabled.
- CEM-13767 Though Contrail fabric manager has the ability for the user to use custom image names for the fabric devices, for platforms like QFX10000-60C which runs on vmhost-based platforms, while uploading the image to CFM, the image name should be chosen in **junos-vmhost-install-x.tgz** format.
- CEM-13685 DPDK vRouter with MLNX CX5 takes about 10 minutes and also lcore crash is seen. This happens once during initial installation.
- CEM-13380 AppFormix Flows does not show up for multi homed devices on the fabric
- CEM-11163 In Fortville X710 NIC: With TX and RX buffers performance degrade is observed as mbufs gets exhausted.
- CEM-10929 When Contrail Insights is querying LLDP table from a device through SNMP, if SNMP calls time out, Contrail Insights marks the device as invalidConfiguration and notifies the user to take a look. When the user verifies that snmpwalk is working and there are no network issues, click **Edit** and reconfigure that device from **Settings > Network Devices** to make Contrail Insights try to run LLDP discovery and add this device again.
- CEM-9979 During upgrade of DPDK computes deployed with OOO Heat Templates in RHOSP environment, vRouter coredumps are observed. This is due to the sequence in which the services are started during upgrade and does not have impact on cluster operation.

- CEM-8701 Onboarding of multiple BMS in parallel on SP-style fabric does not work. While bringing up a BMS using the Life Cycle Management workflow, sometimes on faster servers the re-image does not go through and instance not moved from ironic vn to tenant vn. This is because if the PXE boot request from the BMS is sent before the routes are converged between the BMS port and the TFTP service running in Contrail nodes. As a workaround, the servers can be rebooted or the BIOS in the servers can be configured to have a delayed boot.
- CEM-8149 BMS LCM with fabric set with `enterprise_style=True` is not supported. By default, `enterprise_style` is set to `False`. Avoid using `enterprise_style=True` if the fabric object onboards the BMS LCM instance.
- CEM-5141 For deleting compute nodes, the UI workflow will not work. Instead, update the `instances.yaml` with `"ENABLE_DESTROY: True"` and `"roles:"` (leave it empty) and run the following playbooks.

```
ansible-playbook -i inventory/ -e orchestrator=openstack --tags nova playbooks/
install_openstack.yml
ansible-playbook -i inventory/ -e orchestrator=openstack playbooks/install_contrail.yml
```

For example:

```
global_configuration:
  ENABLE_DESTROY: True
  ...
  ...
instances:
  ...
  ...
  srvr5:
    provider: bms
    ip: 19x.xxx.x.55
    roles:
  ...
  ...
```

- CEM-5043 VNI update on a LR doesn't update the RouteTable. As a workaround, delete the LogicalRouter and create a new LogicalRouter with the new VNI.
- CEM-4370 Additional links cannot be appended to service templates used to create PNF service chaining. If there is a need to add additional links, the service template needs to be deleted and re-added again.

- CEM-4358 In Contrail fabric deployments configuring QFX5110 as spine (CRB-Gateway) does not work.
- CEM-3959 BMS movement across TORs is not supported. To move BMS across TORs the whole VPG needs to be moved. That means if there are more than one BMS associated to one VPG, and one of the BMS need to be moved, the whole VPG need to be deleted and re-configured as per the new association.
- CEM-3245 Multicast traffic originated from type-6 incapable QFX devices are duplicated by vRouters.
- JCB-187287 High Availability provisioning of Kubernetes master is not supported.
- JCB-184776 When the vRouter receives the head fragment of an ICMPv6 packet, the head fragment is immediately enqueued to the assembler. The flow is created as hold flow and then trapped to the agent. If fragments corresponding to this head fragment are already in the assembler or if new fragments arrive immediately after the head fragment, the assembler releases them to flow module. Fragments get enqueued in the hold queue if agent does not write flow action by the time the assembler releases fragments to the flow module. A maximum of three fragments are enqueued in the hold queue at a time. The remaining fragments are dropped from the assembler to the flow module.

As a workaround, the head fragment is enqueued to assembler only after flow action is written by agent. If the flow is already present in non-hold state, it is immediately enqueued to assembler.
- JCB-177787 In DPDK vRouter use cases such as SNAT and LBaaS that require netns, jumbo MTU cannot be set. Maximum MTU allowed: <=1500.

Resolved Issues

You can research limitations that are resolved with this release at:

[Resolved Issues in Contrail Networking Release 2011](#)

Use your Juniper Support login credentials to view the list. If you do not have a Juniper Support account, you can register for one at <https://userregistration.juniper.net/>.

Deprecated Issues

The following feature is deprecated in Contrail Networking Release 2011.L2.

- VMware vCenter

The following features are deprecated in Contrail Networking Release 2008.

- vRouter data plane encryption
- Public cloud vRouter as a multicloud gateway

The following feature is deprecated in Contrail Networking Release 2003.

- Installation and upgrade using Helm deployer